Compressed Absorbing Boundary Conditions for the Helmholtz Equation

by

Rosalie Bélanger-Rioux

B.Sc., McGill University (2009)

Submitted to the Department of Mathematics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2014

© Massachusetts Institute of Technology 2014. All rights reserved.

arXiv:1408.3144v1 [math.NA] 13 Aug 2014

Author .....
Department of Mathematics
May 8, 2014

Certified by .....
Laurent Demanet
Assistant Professor
Thesis Supervisor

Accepted by .....
Michel Goemans
Chairman, Department Committee on Graduate Theses

Compressed Absorbing Boundary Conditions for the Helmholtz Equation

by

Rosalie Bélanger-Rioux

Submitted to the Department of Mathematics
on May 8, 2014, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Absorbing layers are sometimes required to be impractically thick in order to offer an accurate approximation of an absorbing boundary condition for the Helmholtz equation in a heterogeneous medium. It is always possible to reduce an absorbing layer to an operator at the boundary by layer-stripping elimination of the exterior unknowns, but the linear algebra involved is costly. We propose to bypass the elimination procedure, and directly fit the surface-to-surface operator in compressed form from a few exterior Helmholtz solves with random Dirichlet data. We obtain a concise description of the absorbing boundary condition, with a complexity that grows slowly (often, logarithmically) in the frequency parameter. We then obtain a fast (nearly linear in the dimension of the matrix) algorithm for the application of the absorbing boundary condition using partitioned low rank matrices. The result, modulo a precomputation, is a fast and memory-efficient compression scheme of an absorbing boundary condition for the Helmholtz equation.

Thesis Supervisor: Laurent Demanet
Title: Assistant Professor

Acknowledgments

I would like to thank my advisor, Laurent Demanet, for his advice and support throughout my time as a graduate student.

Also, I would like to acknowledge professors I took classes from, including my thesis committee members Steven Johnson and Jacob White, and qualifying examiners Ruben Rosales and again Steven Jonhson. I appreciate your expertise, availability and help.

I would like to thank collaborators and researchers with whom I discussed various ideas over the years, it was a pleasure to learn and discover with you.

Thanks also to my colleagues here at MIT in the mathematics PhD program, including office mates. It was a pleasure to get through this together.

Thanks to my friends, here and abroad, who supported me. From my decision to come to MIT, all the way to graduation, it was good to have you.

Finally, I thank my loved ones for being there for me always, and especially when I needed them.

Contents

1 Introduction **15**
- 1.1 Applications of the Helmholtz equation in an unbounded domain . . . 16
- 1.2 A two-step numerical scheme for compressing ABCs . . . 18
  - 1.2.1 Background material . . . 18
  - 1.2.2 Matrix probing . . . 19
  - 1.2.3 Partitioned low-rank matrices . . . 20
  - 1.2.4 Summary of steps . . . 22

2 Background material **23**
- 2.1 The Helmholtz equation: free-space problem . . . 24
  - 2.1.1 Solution in a uniform medium using the Green's function . . . 24
  - 2.1.2 Solution in a heterogeneous medium using Absorbing Boundary Conditions . . . 25
  - 2.1.3 The Dirichlet-to-Neumann map as an ABC . . . 26
- 2.2 Discrete absorbing boundary conditions . . . 27
  - 2.2.1 Surface-to-surface ABCs . . . 27
  - 2.2.2 Absorbing layer ABCs . . . 28
  - 2.2.3 Complexity of ABCs in heterogeneous media . . . 29
- 2.3 The Helmholtz equation: exterior and half-space problems . . . 30
  - 2.3.1 The exterior problem . . . 30
  - 2.3.2 The half-space problem . . . 31
  - 2.3.3 Solution to the half-space problem in a uniform medium using the Green's function . . . 32
  - 2.3.4 The analytical half-space DtN map . . . 34
- 2.4 Eliminating the unknowns: from any ABC to the DtN map . . . 36
  - 2.4.1 Eliminating the unknowns in the half-space problem . . . 36
  - 2.4.2 Eliminating the unknowns in the exterior problem . . . 38

3 Matrix probing for expanding the Dirichlet-to-Neumann map **41**
- 3.1 Introduction to matrix probing . . . 41
  - 3.1.1 Setup for the exterior problem . . . 42
  - 3.1.2 Matrix probing . . . 43
  - 3.1.3 Solving the Helmholtz equation with a compressed ABC . . . 45
- 3.2 Choice of basis matrices for matrix probing . . . 45
  - 3.2.1 Oscillations and traveltimes for the DtN map . . . 46

3.2.2 More on traveltimes . . . 47
3.3 Numerical experiments . . . 49
3.3.1 Probing tests . . . 52
3.3.2 Using the probed DtN map in a Helmholtz solver . . . 56
3.3.3 Grazing waves . . . 59
3.3.4 Variations of p with N . . . 59
3.4 Convergence of probing for the half-space DtN map: theorem . . . 61
3.4.3 Chebyshev expansion . . . 62
3.4.4 Bound on the derivatives of the DtN map kernel with oscillations removed . . . 63
3.4.5 Bound on the error of approximation . . . 64
3.5 Convergence of probing for the half-space DtN map: numerical confirmation . . . 65
3.5.1 Uniform medium . . . 66

4 Partitioned low-rank matrices for compressing the Dirichlet-to-Neumann map 67
4.1 Partitioned low-rank matrices . . . 68
4.1.1 Construction of a PLR matrix . . . 68
4.1.2 Multiplication of a PLR matrix with a vector . . . 70
4.1.3 Structure and complexity . . . 72
4.2 Using PLR matrices for the DtN map's submatrices . . . 77
4.2.1 Choosing the tolerance . . . 77
4.2.2 Minimizing the matrix-vector application time . . . 78
4.2.3 Numerical results . . . 80
4.3 The half-space DtN map is separable and low rank: theorem . . . 81
4.3.1 Treating the Hankel function . . . 82
4.3.2 Treating the $1/kr$ factor . . . 83
4.3.4 Dyadic interval for the Gaussian quadrature . . . 85
4.3.5 Finalizing the proof . . . 86
4.3.9 The numerical low-rank property of the DtN map kernel for heterogeneous media . . . 89
4.4 The half-space DtN map is separable and low rank: numerical verification 89
4.4.1 Slow disk and vertical fault . . . 90
4.4.2 Constant medium, waveguide, diagonal fault . . . 91
4.4.3 Focusing and defocusing media . . . 92

5 Conclusion 95

A Summary of steps 97
A.1 Matrix probing: $D \to \widetilde{D}$ . . . 98
A.2 PLR compression: $\widetilde{D} \to \overline{D}$ . . . 99
A.3 Using $\overline{D}$ in a Helmholtz solver . . . 99

List of Figures

2-1 The exterior problem: Ω is in grey, there is a thin white strip around Ω, then an absorbing layer in blue. 32
2-2 The half-space problem: there is thin white strip below the bottom side of Ω, and an absorbing layer (in blue) surrounds that strip on three sides. 32

3-1 Color plot of $c(\mathbf{x})$ for the uniform medium. 49
3-2 Color plot of $c(\mathbf{x})$ for the Gaussian waveguide. 49
3-3 Color plot of $c(\mathbf{x})$ for the Gaussian slow disk. 49
3-4 Color plot of $c(\mathbf{x})$ for the vertical fault. 49
3-5 Color plot of $c(\mathbf{x})$ for the diagonal fault. 49
3-6 Color plot of $c(\mathbf{x})$ for the periodic medium. 49
3-7 Condition numbers for the $(1,1)$ block, $c \equiv 1$. 53
3-8 Condition numbers for the $(2,1)$ block, $c \equiv 1$. 53
3-9 Approximation error (line) and probing error (with markers) for the blocks of D, $c \equiv 1$. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$. 54
3-10 Approximation error (line) and probing error (with markers) for the blocks of D, c is the waveguide. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$. 54
3-11 Approximation error for the $(1,1)$ blocks of D, c is the slowness disk, comparing the use of the normal traveltime (circles) to the fast marching traveltime (squares). 55
3-12 Approximation error for the $(2,1)$ blocks of D, c is the slowness disk, comparing the use of the normal traveltime (circles) to the fast marching traveltime (squares). 55
3-13 Approximation error (line) and probing error (with markers) for the blocks of D, c is the fault. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$. 56
3-14 Approximation error (line) and probing error (with markers) for the sub-blocks of the $(1,1)$ block of D, c is the fault. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$. 56
3-15 Real part of the solution, c is the slow disk. 57
3-16 Real part of the solution, c is the waveguide. 57
3-17 Real part of the solution, c is the vertical fault with source on the left. 57
3-18 Real part of the solution, c is the vertical fault with source on the right. 57

3-19 Real part of the solution, c is the diagonal fault. 59
3-20 Imaginary part of the solution, c is the periodic medium. 59
3-21 Error in solution u, $c \equiv 1$, moving point source. Each line in the plot corresponds to using $\widetilde{D}$ from a different row of Table 3.2. 60
3-22 Probing error of the $(1,1)$ block of the DtN map for the slow disk, fixed FD error level of 10^{-1}, increasing N. This is the worst case, where p follows a weak power law with N. 60
3-23 Probing error of the $(1,1)$ block of the DtN map for the waveguide, fixed FD error level of 10^{-1}, increasing N. Here p follows a logarithmic law. 60
3-24 Probing error of the half-space DtN map($q = 1$, 10 trials, circle markers and error bars) compared to the approximation error (line), $c \equiv 1$, $L = 1/4$, $\alpha = 2$, $n = 1024$, $\omega = 51.2$. 65
3-25 Condition numbers for probing the half-space DtN map, $c \equiv 1$, $L = 1/4$, $\alpha = 2$, $n = 1024$, $\omega = 51.2$, $q = 1$, 10 trials. 65

4-1 Weak hierarchical structure, $\frac{N}{R_{\max}} = 8$. 74
4-2 Strong hierarchical structure, $\frac{N}{R_{\max}} = 16$. 74
4-3 Corner PLR structure, $\frac{N}{R_{\max}} = 8$. 74
4-4 Matrix-vector complexity for submatrix $(1,1)$ for $c \equiv 1$, various $R_{\max}$. Probing errors of 10^{-5}, 10^{-6}. 78
4-5 Matrix-vector complexity for submatrix $(2,1)$ for $c \equiv 1$, various $R_{\max}$. Probing errors of 10^{-5}, 10^{-6}. 78
4-6 PLR structure of the probed submatrix $(1,1)$ for $c \equiv 1$, $R_{\max} = 2$. Each block is colored by its numerical rank. 79
4-7 PLR structure of the probed submatrix $(1,1)$ for $c \equiv 1$, $R_{\max} = 8$. Each block is colored by its numerical rank. 79
4-8 PLR structure of the probed submatrix $(1,1)$ for $c \equiv 1$, $R_{\max} = 32$. Each block is colored by its numerical rank. 79
4-9 PLR structure of the probed submatrix $(2,1)$ for $c \equiv 1$, $R_{\max} = 4$. Each block is colored by its numerical rank. 79
4-10 Maximum off-diagonal ranks with N for the slow disk, various ϵ. FD error of 10^{-3}, $r_0 = h$. 90
4-11 Maximum off-diagonal ranks with N for the slow disk, various ϵ. FD error of 10^{-3}, $r_0 = 4h$. 90
4-12 Maximum off-diagonal ranks with ε for the slow disk, various N. FD error of 10^{-3}, $r_0 = h$. 91
4-13 Maximum off-diagonal ranks with ε for the slow disk, various N. FD error of 10^{-3}, $r_0 = 4h$. 91
4-14 Maximum off-diagonal ranks with N for the diagonal fault, various ϵ. FD error of 10^{-2}, $r_0 = h$. 91
4-15 Maximum off-diagonal ranks with ε for the diagonal fault, various N. FD error of 10^{-2}, $r_0 = h$. 91

4-16 Maximum off-diagonal ranks for the diagonal fault as a function of N, various tolerances ϵ. The separation is $r_0 = h$. FD error of 10^{-3}. . . . 92
4-17 Maximum off-diagonal ranks for the diagonal fault as a function of the tolerance ϵ, various N. The separation is $r_0 = h$. FD error of 10^{-3}. . 92
4-18 Maximum off-diagonal ranks for the diagonal fault as a function of N, various tolerances ϵ. The separation is $r_0 = 8h$. FD error of 10^{-2}. . . 92
4-19 Maximum off-diagonal ranks for the diagonal fault as a function of the tolerance ϵ, various N. The separation is $r_0 = 8h$. FD error of 10^{-2}. . 92
4-20 Maximum off-diagonal ranks for the focusing medium (4.21) as a function of N, various tolerances ϵ. Separation is $r_0 = h$, FD error of 10^{-2}. 93
4-21 Maximum off-diagonal ranks for the focusing medium (4.21) as a function of the tolerance ϵ, various N. Separation is $r_0 = h$, FD error of 10^{-2}. 93
4-22 Maximum off-diagonal ranks for the focusing medium (4.21) as a function of N, various tolerances ϵ. Separation is $r_0 = h$, FD error of 10^{-3}. 94
4-23 Maximum off-diagonal ranks for the focusing medium (4.21) as a function of the tolerance ϵ, various N. Separation is $r_0 = h$, FD error of 10^{-3}. 94

List of Tables

3.1 For each medium considered, we show the parameters N and $\omega/2\pi$, along with the resulting discretization error caused by the Finite Difference (FD error) formulation. We also show the width w of the pPML needed, in number of points, to obtain an error caused by the pPML of less than $1e-1$. Furthermore, we show the total number P of basis matrices needed to probe the entire DtN map with an accuracy of about $1e-1$ as found in Section 3.3.1. Finally, we show the position of the point source used in calculating the solution u. 50
3.2 $c \equiv 1$ 58
3.3 c is the waveguide 58
3.4 c is the slow disk 58
3.5 c is the fault 58
3.6 c is the diagonal fault 58
3.7 c is the periodic medium 59

4.1 PLR compression results, $c \equiv 1$ 80
4.2 PLR compression results, c is the Gaussian waveguide. 80
4.3 PLR compression results, c is the Gaussian slow disk. 81
4.4 PLR compression results, c is the vertical fault, sources on the left and on the right. 81
4.5 PLR compression results, c is the diagonal fault. 82
4.6 PLR compression results, c is the periodic medium. 82

Chapter 1

Introduction

This work investigates arbitrarily accurate realizations of absorbing (a.k.a. open, radiating) boundary conditions (ABC), including absorbing layers, for the 2D acoustic high-frequency Helmholtz equation in certain kinds of heterogeneous media. Instead of considering a specific modification of the partial differential equation, such as a perfectly matched layer, we study the broader question of compressibility of the nonlocal kernel that appears in the exact boundary integral form $Du = \partial_\nu u$ of the ABC. The operator D is called the *Dirichlet-to-Neumann* (DtN) map. This boundary integral viewpoint invites to rethink ABCs as a two-step numerical scheme, where

1. a precomputation sets up an expansion of the kernel of the boundary integral equation, then

2. a fast algorithm is used for each application of this integral kernel at the open boundaries in a Helmholtz solver.

This two-step approach may pay off in scenarios when the precomputation is amortized over a large number of solves of the original equation with different data. This framework is, interestingly, half-way between a purely analytical or physical method and a purely numerical one. It uses both the theoretical grounding of analytic knowledge and the intuition from understanding the physics of the problem in order to obtain a useful solution.

The numerical realization of ABC typically involves absorbing layers that become impractical for difficult $c(\mathbf{x})$, or for high accuracy. We instead propose to realize the ABC by directly compressing the integral kernel of D, so that the computational cost of its setup and application would become competitive when (2.1) is to be solved multiple times. Hence this paper is not concerned with the design of a new ABC, but rather with the reformulation of existing ABCs that otherwise require a lot of computational work per solve. In many situations of practical interest we show that it is possible to "learn" the integral form of D, as a precomputation, from a small number of solves of the exterior problem with the expensive ABC. By "small number", we mean a quantity essentially independent of the number of discretization points N along one dimension – in practice as small as 1 or as large as 50. We call this strategy

matrix probing. To show matrix probing is efficient, we prove a result on approximating D in a special case, with inverse powers multiplied by a complex exponential. This leads us to the successful design of a basis for a variety of heterogeneous media.

Once we obtain a matrix realization $\widetilde{D}$ of the ABC from matrix probing, we can use it in a Helmholtz solver. However, a solver would use matrix-vector multiplications to apply the dense matrix $\widetilde{D}$. Hence the second step of our numerical scheme: we compress $\widetilde{D}$ using partitioned low-rank (PLR) matrices to acquire a fast matrix-vector product. This second step can only come after the first, since it is the first step that gives us access to the entries in D and allows us to use the compression algorithms of interest to us. We know we can use hierarchical or partitioned-low-rank matrices to compress the DtN map because we prove the numerical low-rank property of off-diagonal blocks of the DtN map, and those algorithms exploit favorably low-rank blocks. Since PLR matrices are more flexible than hierarchical matrices, we use them to compress $\widetilde{D}$ into $\overline{D}$ and apply it to vectors in complexity ranging from $O(N \log N)$ (more typical) to $O(N^{3/2})$ (worst case). The precomputation necessary to set up the PLR approximation is of similar complexity. This can be compared to the complexity of a dense matrix-vector product, which is $O(N^2)$.

In this introduction, we first motivate in section 1.1 the study of the Helmholtz equation in an unbounded domain by presenting important applications. We then give more details on the steps of our numerical scheme in section 1.2, which will make explicit the structure of this thesis.

1.1 Applications of the Helmholtz equation in an unbounded domain

We consider the scalar Helmholtz equation in $\mathbb{R}^2$,

$$\Delta u(\mathbf{x}) + \frac{\omega^2}{c^2(\mathbf{x})} u(\mathbf{x}) = f(\mathbf{x}), \qquad \mathbf{x} = (x_1, x_2), \tag{1.1}$$

with compactly supported f, the *source*. We seek the solution u. The function c in (1.1) is called the *medium*, ω the frequency.

We consider the unique solution u to (1.1) determined by the Sommerfeld radiation condition (SRC), which demands that the solution be outgoing. We call the problem of finding a solution u to (1.1) with the SRC the *free-space problem*. There are many applications that require a solution to the free-space Helmholtz problem, or the related free-space Maxwell's equations problem. To solve the free-space problem numerically, we reformulate the problem to a bounded domain Ω in which we shall find the solution u. One must then impose an *absorbing boundary condition* (ABC) on the boundary $\partial\Omega$. ABCs are designed to absorb waves impinging on $\partial\Omega$ so the waves do not reflect back in Ω and pollute the solution there. ABCs in heterogeneous media are often too costly, but the two-step scheme of this thesis addresses this problem. Let us now discuss applications of this two-step scheme.

The main application is wave-based imaging, an inverse problem for the wave

equation

$$\Delta U(\mathbf{x},t) - \frac{1}{c^2(\mathbf{x})}\frac{\partial^2 U(\mathbf{x},t)}{\partial t^2} = F(\mathbf{x},t), \tag{1.2}$$

and related equations. The Helmholtz equation is equivalent to the wave equation because we can decompose the solution U and the source F into time-harmonic components by a Fourier transform in time. Another type of wave equation, the Maxwell's equations, can also be reformulated as Helmholtz equations on the various components of the electric and magnetic fields.

An inverse problem is as follows: instead of trying to find the solution u of (1.1) given ω, f and c, we do the opposite. In other words, we are given the solution u at a set of receiver locations for some ω and various sources f. We try to determine the medium c from that information. We are usually interested in (or can only afford) $c(\mathbf{x})$ for $\mathbf{x}$ in some part of the whole space, say some bounded domain Ω, and absorbing boundary conditions are then necessary. To solve the inverse problem in Ω, we need to use a lot of sources, say a thousand. The details of why many sources are useful, and how to solve the inverse problem, are not relevant to this thesis. What is relevant is that, in the course of solving the inverse problem for say the Helmholtz equation, it is needed to solve the free-space Helmholtz problem for all these sources, with ABCs, in heterogeneous media. We list here some imaging applications where the Helmholtz equation, or other types of wave equations, are used to solve inverse problems with ABCs in heterogeneous media, that is, where the numerical scheme in this thesis might prove useful.

- In seismic imaging, we acquire knowledge of the rock formations under the earth's surface. That is, we want to know the medium c in which the waves propagate. The domain Ω might be a region surrounding a seismic fault [54] where one wants to assess earthquake hazards, or a place where one might like to find hydrocarbons, other minerals, or even geothermal energy [35]. ABCs are needed to simulate the much larger domain in which Ω is embedded, that is, the Earth. Sources (which might be acoustic or electromagnetic) and receivers might be placed on the surface of the Earth or inside a well, or might be towed by a boat or placed at the bottom of the sea. An earthquake can also be used as a source.

- Ultrasonic testing [28] is a form of non-destructive testing where very short ultrasonic pulses are sent inside an object. The object might be, say, a pipe that is being tested for cracks or damage from rust, or a weld being tested for defects. The received reflections or refractions from those ultrasonic pulses are used for diagnosis on the object of interest. Ω might be the object itself, or a part of it. Ultrasonic imaging [52] is also used in medicine to visualize fetuses, muscles, tendons and organs. The domain Ω is then the body part of interest.

- Synthetic-aperture radar imaging is used to visualize a scene by sending electromagnetic pulses from an antenna aboard a plane or satellite [10]. It is also used to detect the presence of an object [7] far away such as a planet, or through clouds or foliage.

An entirely different application of the Helmholtz equation resides in photonics [33]. The field of photonics studies the optical properties of materials. In particular, one tries to construct photonic crystals, a periodic medium with desired properties depending on the specific application. It is of course useful to first test photonic crystals numerically to observe their properties before actually building them. However, since crystals are assumed to be infinite, absorbing boundary conditions need to be used to reformulate the problem on a bounded domain. This is where our two-step numerical scheme can be used.

1.2 A two-step numerical scheme for compressing ABCs

The next chapter, chapter 2, will present theoretical facts about the Helmholtz equation and related concepts, which will be useful at various points in later chapters for developing the proposed framework. Then, chapter 3 will present the first step of the scheme: matrix probing and how it is used for providing a rapidly converging expansion of the DtN map. Next, chapter 4 will contain the second part of the proposed two-step scheme: the compression of the matrix probing expansion, for fast evaluation, using partitioned low-rank matrices. After the precomputations of matrix probing and PLR matrices, we finally obtain a fast and accurate compressed absorbing boundary condition. Chapter 5 concludes this thesis with a review of the main new ideas presented, identification of further research directions and open problems, and also an overview of how the presented framework could be used in other contexts. A summary of the steps involved in the presented two-step numerical scheme is available in appendix A.

1.2.1 Background material

We begin by introducing again the problem we wish to solve, the *free-space problem* for the Helmholtz equation. This problem is defined on an unbounded domain such as $\mathbb{R}^2$, but can be solved by reformulating to a bounded domain we will call Ω. One way to do this is to impose what we call the *exterior Dirichlet-to-Neumann map* (DtN map) on $\partial\Omega$. The DtN map D relates Dirichlet data u on $\partial\Omega$ to the normal derivative $\partial_\nu u$ of u, where ν is the unit outward vector to $\partial\Omega$: $Du = \partial_\nu u$. This allows us to solve the Helmholtz equation on Ω to obtain a solution u which coincides, inside Ω, with the solution to the free-space problem.

We shall see that any *absorbing boundary condition* (ABC) or Absorbing Layer (AL) can be used to obtain the DtN map. As we mentioned before, an ABC is a special condition on the boundary $\partial\Omega$ of Ω, which should minimize reflections of waves reaching $\partial\Omega$. Different ABCs work differently however, and so we will review existing techniques for constructing ABCs for the Helmholtz equation, and see again how they are computationally expensive in a variable medium. This suggests finding a different, more efficient, way of obtaining the exterior DtN map from an ABC. To attack this problem, we consider the *half-space* DtN map, known analytically in a

constant medium. This half-space DtN map is actually quite similar to the exterior DtN map, at least when Ω is a rectangular domain. Indeed, the interface along which the DtN map is defined for the half-space problem is a straight infinite line $x_2 = 0$. For the exterior problem on a rectangle, the DtN map is defined on each of the straight edges of Ω. Hence the restriction of the exterior DtN map to one such edge of Ω, which say happens to be on $x_2 = 0$, behaves similarly to the restriction of the half-space DtN map to that same edge. The main difference between those two maps is created by scattering from corners of $\partial\Omega$. Both in chapter 3 and 4, we will prove facts about the half-space DtN map which will inform our numerical scheme for the exterior DtN map.

We end the chapter of background material with how the most straightforward way of obtaining the exterior DtN map from an ABC, *layer-stripping*, is prohibitively slow, especially in variable media. This will also explain how, even if we have an efficient way of obtaining the DtN map from an ABC, applying it at every solve of the Helmholtz equation will be slow. In fact, once we have the map D in $Du = \partial_\nu u$, we need to apply this map to vectors inside a Helmholtz solver. But D has dimension approximately N, if N is the number of points along each direction in Ω, so matrix-vector products with D have complexity N^2. This is why we have developed the two-step procedure presented in this thesis: first an expansion of the DtN map using matrix probing, then a fast algorithm using partitioned low-rank matrices for the application of the DtN map.

1.2.2 Matrix probing

Chapter 3 is concerned with the first step of this procedure, namely, setting up an expansion of the exterior DtN map kernel, in a precomputation. This will pave the way for compression in step two, presented in chapter 4.

Matrix probing is used to find an expansion of a matrix M. For this, we assume an expansion of the type

$$M \approx \widetilde{M} = \sum_{j=i}^{p} c_j B_j,$$

where the *basis matrices* $\{B_j\}$ are known, and we wish to find the coefficients $\{c_j\}$. We do not have access to M itself, but only to products of M with vectors. In particular, we can multiply M with a random vector z to obtain

$$w = Mz \approx \sum_{j=1}^{p} c_j B_j z = \Psi_z \, \mathbf{c}.$$

We can thus obtain the vector of coefficients $\mathbf{c}$ by applying the pseudoinverse of Ψ_z to $Mz = w$.

For matrix probing to be an efficient expansion scheme, we need to carefully choose the basis matrices $\{B_j\}$. Here, we use knowledge of the half-space DtN map to inform our choices: a result on approximating the half-space DtN map with a particular set of functions, inverse powers multiplied by a complex exponential. Hence our basis

matrices are typically a discretization of the kernels

$$B_j(x,y) = \frac{e^{ik|x-y|}}{(h+|x-y|)^{j/2}}, \tag{1.3}$$

where h is our discretization parameter, $h = 1/N$. The need for a careful design of the basis matrices can however be a limitation of matrix probing. In this thesis, we have used also insights from geometrical optics to derive basis matrices which provide good convergence in a variety of cases. Nonetheless, in a periodic medium such as a photonic crystal, where the wavelength is as large as the features of the medium, the geometrical optics approximation breaks down. Instead, we use insights from the solution in a periodic medium, which we know behaves like a Bloch wave, to design basis matrices. However, the periodic factor of a Bloch wave does not lead to very efficient basis matrices since it is easily corrupted by numerical error. Another limitation is that a medium which has discontinuities created discontinuities as well in the DtN map, forcing again a more careful design of basis matrices.

Once we do have basis matrices, we can probe the DtN map D, block by block. Indeed, we choose $\Omega = [0,1]^2$, and so $\partial\Omega$ is the four edges of the square[1], numbered counter-clockwise. Hence, D has a 4 by 4 structure corresponding to restrictions of the DtN map to each pair of edges. To obtain the product of the (i_M, j_M) block M of D with a random vector, we need to solve what we call the *exterior problem*: we put a random Dirichlet boundary condition on edge j_M of $\partial\Omega$, solve the Helmholtz equation outside of Ω using an ABC, and take the normal derivative of the solution on edge i_M of $\partial\Omega$.

In situations of practical interest, we obtain the DtN map as a precomputation using matrix probing with leading complexity of about 1 to 50 solves of the exterior problem with the expensive ABC – a number of solves essentially independent of the number of discretization points N. A solve of the exterior problem is essentially equivalent to a solve of the original problem with that same expensive ABC.

We then present a careful study of using matrix probing to expand the DtN map in various media, use the matrix probing expansion to solve the Helmholtz equation, and document the complexity of the method.

In the next chapter, we will present a fast algorithm for applying the DtN map's expansion found by matrix probing.

1.2.3 Partitioned low-rank matrices

In chapter 4, we produce a fast algorithm for applying the DtN map D to vectors, since this is an operation that a Helmholtz solver needs. Again, we needed matrix probing first, to obtain an explicit approximation of D, to be compressed in order to obtain a fast matrix-vector product. Indeed, we do not have direct access to the entries of D at first, but rather we need to solve a costly problem, the exterior problem, every time we need a multiplication of D with a vector. Now that we have an

[1]The framework does carry over to polygonal domains easily, but we do not cover this here.

explicit representation of D from matrix probing, with cost $O(N^2)$ for matrix-vector products, we can compress that representation to obtain a faster product.

We have mentioned before how the half-space DtN map is similar to the exterior DtN map. In chapter 4, we will prove that the half-space DtN map kernel K in constant medium is numerically low rank. This means that, away from the singularity of K, the function $K(|x-y|)$ can be written as a short sum of functions of x multiplied by functions of y:

$$K(|x-y|) = \sum_{j=1}^{J} \Psi_j(x)\chi_j(y) + E(x,y)$$

with error E small. The number of terms J depends logarithmically on the error tolerance and on the frequency k. This behavior carries over to some extent to the exterior DtN map. Thus we use a compression algorithm which can exploit the low-rank properties of blocks of D that are not on the diagonal, that is, that are away from the singularity.

A well-known such compression framework is called *hierarchical matrices*. Briefly, hierarchical matrices adaptively divide or compress diagonal blocks. We start by dividing the matrix in 4 blocks of half the original matrix's dimension. The two off-diagonal blocks are compressed: we express them by their singular value decomposition (SVD), truncated after R terms. The two diagonal blocks are divided in four again, and we recurse: off-diagonal blocks are compressed, diagonal blocks are divided, etc. We do not divide a block if compressing it will result in less error than our error tolerance ε – this is adaptivity. The parameters R and ε are chosen judiciously to provide a fast matrix-vector multiplication with small error.

However, the hierarchical matrix framework can only apply to matrices with a singularity along the diagonal. This is not useful to us, since for example a block of D corresponding to two consecutive edges of $\partial\Omega$ will have the singularity in a corner. We thus decide to use partitioned low-rank (PLR) matrices for compressing and applying D. PLR matrices are more flexible than hierarchical matrices: when we divide a block, or the original matrix, in 4 sub-blocks, any of those four sub-blocks can be divided again. Any block that is not divided any further is called a *leaf*, and $\mathcal{B}$ is the set of all leaves. If a singularity is in a corner, then the PLR compression algorithm will automatically divide blocks close to that corner, but will compress farther blocks since they have lower numerical rank. Note that we use the randomized SVD [39] to speed us the compression, so that its complexity is of order $O(NR^2|\mathcal{B}|)$, where $|\mathcal{B}|$ is often on the order of $\log N$ but will be $\sqrt{N}$ in the worst case. Similarly, the complexity of a matrix-vector product is usually $O(NR|\mathcal{B}|)$, which for $N \approx 1000$ provides a speed-up over a dense matrix-vector product of a factor of 30 to 100. We also show that the worst case complexity of a matrix-vector product in the PLR framework is $O(N^{3/2})$. This should be compared to the complexity of a dense matrix-vector product, which is $O(N^2)$.

We may then use PLR matrices to compress the DtN map, and use this compressed map in a Helmholtz solver. We verify the complexity of the method, and present results on the solution of the Helmholtz equation with a probed and compressed DtN map.

1.2.4 Summary of steps

In appendix A, we present a summary of the various operations involved in implementing the presented two-step algorithm.

Chapter 2

Background material

In this chapter, we review the necessary background material for this thesis. We begin by introducing in section 2.1 the problem we wish to solve, that is, the *free-space problem* for the Helmholtz equation. This problem is defined on an unbounded domain. We will explain how *absorbing boundary conditions* (ABCs) and the *exterior Dirichlet-to-Neumann map* (DtN map) are related, and can be used to reformulate the free-space problem from an unbounded domain to a bounded one. This reformulation allows us to solve the free-space problem.

We then briefly review existing techniques for constructing ABCs for the Helmholtz equation in section 2.2. From this, we will understand why, when solving the Helmholtz equation in a heterogeneous medium, ABCs are computationally expensive. This suggests an approach where we compress an ABC, or the related exterior DtN map. In other words, we find a way to apply the ABC which allows for a faster solve of the free-space problem.

To attack this problem, we first look at the *half-space* DtN map, which is known analytically in uniform medium. In some respect, this half-space DtN map is quite similar to the exterior DtN map. To understand this similarity, we introduce in section 2.3 the concepts of the *exterior* and *half-space* problems for the Helmholtz equation, and the related half-space DtN map. In the next chapters, we will prove facts about the half-space DtN map which will inform our compression scheme for the exterior DtN map.

Before we explain our work, we end this chapter of background material with section 2.4. We first describe how we may eliminate unknowns in the Helmholtz discretized system to obtain a Riccati equation governing the half-space DtN map. We then describe how the most straightforward way of compressing the exterior DtN map, by eliminating the exterior unknowns in the Helmholtz discretized system, which we call *layer-stripping*, is prohibitively slow. This will explain also how, even if we have a more efficient way of obtaining the DtN map, applying it at every solve of the Helmholtz equation might still be slow. This is why we have developed the two-step procedure presented in this thesis: first an expansion of the DtN map, then a fast algorithm for the application of the DtN map.

2.1 The Helmholtz equation: free-space problem

We consider the scalar Helmholtz equation in $\mathbb{R}^2$,

$$\Delta u(\mathbf{x}) + \frac{\omega^2}{c^2(\mathbf{x})} u(\mathbf{x}) = f(\mathbf{x}), \qquad \mathbf{x} = (x_1, x_2), \tag{2.1}$$

with compactly supported f. We shall call this f the *right-hand side* or the *source*. Here, the solution we seek to this equation is u. The function c in (2.1) is called the *medium of propagation*, or simply *medium*. When c is a constant, we say the medium is *uniform*. When c varies, we say the medium is *heterogeneous*. We call ω the frequency, and note (this shall be explained later) that a high frequency makes the problem harder to solve numerically.

Throughout this work we consider the unique solution u to (2.1) determined by the Sommerfeld radiation condition (SRC) at infinity: when $c(\mathbf{x})$ extends to a uniform c outside of a bounded set[1], the SRC is [43]

$$\lim_{r\to\infty} r^{1/2} \left(\frac{\partial u}{\partial r} - iku \right) = 0, \qquad k = \frac{\omega}{c}, \tag{2.2}$$

where r is the radial coordinate. We call the problem of finding a solution u to (2.1) and (2.2) the *free-space problem*.

2.1.1 Solution in a uniform medium using the Green's function

When the medium c is uniform, there exists an analytical solution to the free-space problem, using the *Green's function*.

Definition 1. *The free-space* Green's function *for the Helmholtz equation* (2.1) *is the unique function* $G(\mathbf{x}, \mathbf{y})$ *which solves the free-space problem with right-hand side a delta function,* $f(\mathbf{x}) = \delta(|\mathbf{x} - \mathbf{y}|)$*, for every fixed* $\mathbf{y}$*.*

It is well-known (p. 19 of [11]) that the Green's function G of the uniform free-space problem (2.1), (2.2) is the following:

$$G(\mathbf{x}, \mathbf{y}) = \frac{i}{4} H_0^{(1)}(k|\mathbf{x} - \mathbf{y}|) \tag{2.3}$$

where $H_0^{(1)}$ is the Hankel function of zeroth order of the first kind. Then, one can compute the solution $u(\mathbf{x})$ for any $\mathbf{x} \in \mathbf{R}^2$, with now any given right-hand side f

[1] If the medium $c(\mathbf{x})$ does not extend to a uniform value outside of a bounded set, it is possible one could use the limiting absorption principle to define uniqueness of the solution to the Helmholtz equation. Start from the wave equation solution $u(\mathbf{x}, t)$. We would like to take the Fourier transform in time of u to obtain the Helmholtz solution. We may take the Laplace transform of u instead to obtain $\widehat{u}(\mathbf{x}, s)$, and ask that the path of integration in the s variable approach the imaginary axis from the decaying side. This could then be used to define the Helmholtz solution.

supported on some bounded domain, using the following formula (see p. 62 of [24]):

$$u(\mathbf{x}) = \int_{\mathbf{R}^2} G(|\mathbf{x} - \mathbf{y}|) f(\mathbf{y}) \, dy. \tag{2.4}$$

To rapidly and accurately evaluate the previous expression at many $\mathbf{x}$'s is not easy: $G(\mathbf{x}, \mathbf{y})$ has a singularity at $\mathbf{x} = \mathbf{y}$, and so care must be taken in the numerical evaluation of the integral in (2.4). Since the issue of evaluating the solution to the free-space problem when the Green's function is known is not directly relevant to this thesis, we refer the reader to [45].

2.1.2 Solution in a heterogeneous medium using Absorbing Boundary Conditions

For a number of practical applications of the Helmholtz equation, as was mentioned in the introductory chapter, one cannot find an analytical solution to this problem because the medium is not uniform, that is, c is not a constant. In particular, the Green's function is not known, and one expects that calculating this Green's function numerically is just as hard as calculating the solution with numerical methods.

To obtain a numerical solution to an equation on an unbounded domain then, one must first reformulate this equation to obtain a possibly modified equation on a bounded domain. Hence, we pick a domain Ω in $\mathbb{R}^2$, with the compact support of f contained in Ω, such that Ω contains the area of interest, that is, where we care to obtain a solution. We now seek to reformulate the SRC on the boundary $\partial\Omega$, so that the resulting solution inside Ω matches that of the free-space problem. This leads us to define an *Absorbing Boundary Condition*:

Definition 2. *An* Absorbing Boundary Condition *for the Helmholtz equation* (2.1) *is a condition on $\partial\Omega$, the boundary of a closed, bounded domain $\Omega \in \mathbb{R}^2$, which uniquely defines a solution to the Helmholtz equation restricted to Ω, such that this unique solution matches the solution to the free-space problem* (2.1), (2.2).

Clearly, if we can reformulate the SRC on the boundary $\partial\Omega$, so that the resulting solution inside Ω matches that of the free-space problem, we will obtain an ABC.

ABCs are extremely important to the numerical solution of the free-space problem because, as alluded to previously, they allow us to restrict the computational domain to a bounded domain Ω, where a solution can be computed in finite time with finite memory. We will discuss ABCs in more detail in the next section, section 2.2.

We will explain in the next chapter the particular, quite rudimentary, solver we used in our numerical experiments. We note here that a better solver should be used for treating larger problems or obtaining better accuracy, for two reasons. First of all, as we explain in more detail in section 2.4, the cost of solving the Helmholtz problem with a standard solver in two dimensions is $O(N^4)$, which is prohibitive. Secondly, as we discuss in section 2.2.3, a higher frequency means we need more points per wavelength in our discretization – this is known as the *pollution effect*. To treat larger problems, there exist better solvers such as the sweeping preconditioner of

Engquist and Ying [20, 19], the shifted Laplacian preconditioner of Erlangga [22, 21], the domain decomposition method of Stolk [49], or the direct solver with spectral collocation of Martinsson, Gillman and Barnett [26, 25]. The problem of appropriate numerical solvers for the Helmholtz equation in high frequency is very much a subject of ongoing research and not the purpose of this thesis, hence we do not discuss this further.

2.1.3 The Dirichlet-to-Neumann map as an ABC

We now seek to reformulate the SRC (2.2) on $\partial\Omega$. There are many ways to do that numerically, as we shall see in section 2.2, but we wish to highlight this particular, analytical way because it introduces a fundamental concept, the *Dirichlet-to-Neumann* map.

Let $G(\mathbf{x}, \mathbf{y})$ be the Green's function for the free-space problem. Define the single and double layer potentials, respectively, on some closed contour Γ by the following, for $\psi,\ \phi$ on Γ (see details in [43], [13]):

$$S\psi(\mathbf{x}) = \int_\Gamma G(\mathbf{x}, \mathbf{y})\ \psi(\mathbf{y})\ dS_y, \qquad T\phi(\mathbf{x}) = \int_\Gamma \frac{\partial G}{\partial \nu_{\mathbf{y}}}(\mathbf{x}, \mathbf{y})\ \phi(\mathbf{y})\ dS_{\mathbf{y}},$$

where ν is the outward pointing normal to the curve Γ, and $\mathbf{x}$ is not on Γ. Now let u^+ satisfy the Helmholtz equation (2.1) in the exterior domain $\mathbb{R}\setminus\overline{\Omega}$, along with the SRC (2.2). Then Green's third identity is satisfied in the exterior domain: using $\Gamma = \partial\Omega$, we get

$$Tu^+ - S\frac{\partial u}{\partial \nu}^+ = u^+, \qquad \mathbf{x} \in \mathbb{R}^2 \setminus \overline{\Omega}. \tag{2.5}$$

Finally, using the jump condition of the the double layer T, we obtain Green's identity on the boundary $\partial\Omega$:

$$(T - \frac{1}{2}I)\, u^+ - S\frac{\partial u}{\partial \nu}^+ = 0, \qquad \mathbf{x} \in \partial\Omega.$$

When the single-layer potential S is invertible[2], we can let $D = S^{-1}(T - \frac{1}{2}I)$, and equivalently write (dropping the + in the notation)

$$\frac{\partial u}{\partial \nu} = Du, \qquad \mathbf{x} \in \partial\Omega. \tag{2.6}$$

The operator D is called the *exterior Dirichlet-to-Neumann map* (or DtN map), because it maps the Dirichlet data u to the Neumann data $\partial u/\partial\nu$ with ν pointing outward. The DtN map is independent of the right-hand side f of (2.1) as long as f is supported in Ω. The notion that (2.6) can serve as an exact ABC was made clear in a uniform medium, e.g., in [18] and in [37]. Equation (2.6) continues to hold

[2]This is the case when there is no interior resonance at frequency ω, which could be circumvented by the use of combined field integral equations as in [13]. The existence and regularity of D ultimately do not depend on the invertibility of S.

even when $c(\mathbf{x})$ is heterogenous in the vicinity of $\partial\Omega$, provided the correct (often unknown) Green's function is used. The medium is indeed heterogeneous near $\partial\Omega$ in many situations of practical interest, such as in geophysics.

The DtN map D is symmetric. The proof of the symmetry of D was shown in a slightly different setting here [38] and can be adapted to our situation. Much more is known about DtN maps, such as the many boundedness and coercivity theorems between adequate fractional Sobolev spaces (mostly in free space, with various smoothness assumptions on the boundary). We did not attempt to leverage these properties of D in the scheme presented here.

We only compress the exterior DtN map in this work, and often refer to it as the DtN map for simplicity, unless there could be confusion with another concept, for example with the *half-space DtN map*. We shall talk more about the half-space DtN map soon, but first, we review in the upcoming section existing methods for discrete absorbing boundary conditions.

2.2 Discrete absorbing boundary conditions

There are many ways to realize an absorbing boundary condition for the Helmholtz equation, and we briefly describe the main ones in this section. We start with ABCs that are surface-to-surface, and move on to ABCs which involve surrounding the computational domain Ω by an absorbing layer. The later approach is often more desirable because the parameters of the layer can usually be adjusted to obtain a desired accuracy. We then discuss the complexity of ABCs in heterogeneous media.

2.2.1 Surface-to-surface ABCs

An early seminal work in absorbing boundary conditions was from Engquist and Majda, who in [18] consider the half-space problem ($x_1 \geq 0$) for the wave equation, with uniform medium, and write down the form of a general wave packet traveling to the left (towards the negative x_1 direction). From this, they calculate the boundary condition which exactly annihilates those wave packets, and obtain [3]:

$$d/dx - i\sqrt{\omega^2 - \xi^2} \tag{2.7}$$

where (ω, ξ) are the dual variables to (y, t) in Fourier space. They can then approximate the square root in various ways in order to obtain usable, i.e. local in both space and time, boundary conditions, recalling that $i\omega$ corresponds to $\partial/\partial y$ and $i\xi$ corresponds to $\partial/\partial t$.

Hagstrom and Warburton in [31] also consider the half-space problem, take the transverse Fourier-Laplace transforms of the solution and use a given Dirichlet data on the boundary of the half-space to obtain what they call a complete wave representation of the solution, valid away from the boundary. They then use this representation to

[3] We are omitting details here for brevity, including a discussion of pseudo-differential operators and their symbols.

obtain approximate local boundary condition sequences to be used as ABC's. Again, their method was developed for the uniform case.

Keller and Givoli, in [37], use a different technique: they assume a circular or spherical Ω, and a uniform medium outside of this Ω. They can then use the Green's function of the exterior problem, which is known for a circle, in order to know the solution anywhere outside Ω, given the boundary data u on $\partial\Omega$. They can then differentiate this solution in the radial (which is the normal) coordinate, and evaluate it on $\partial\Omega$ to obtain the exterior DtN map. They can now use this DtN map as an ABC in a Helmholtz solver. This technique can be seen as *eliminating the exterior unknowns*: as we do not care to know the solution outside of Ω, we can use the information we have on the exterior solution to reduce the system to one only on the inside of Ω. This meaning of *eliminating the exterior unknowns* shall become more obvious when we apply this to the discretized Helmholtz equation in section 2.4.

Somersalo et al. ([47]) also use the DtN map, this time to solve an interior problem related to the Helmholtz equation. They use a differential equation of Riccati type to produce the DtN map. When introduce the elimination of unknowns in section 2.4, we shall demonstrate the connection we have found between the *half-space* problem and a Riccati equation for the DtN map.

The aforementioned ABCs are not meant to be a representative sample of all the existing techniques. Unfortunately, these techniques either do not apply to heterogeneous medium, or do not perform very well in that situation. In contrast, various types of absorbing layers can be used in heterogeneous media, with caution, and are more flexible.

2.2.2 Absorbing layer ABCs

Another approach to ABCs is to surround the domain of interest by an *absorbing layer*. While a layer should preferably be as thin as possible, to reduce computational complexity, its design involves at least two different factors: 1) waves that enter the layer must be significantly damped before they re-enter the computational domain, and 2) reflections created when waves cross the domain-layer interface must be minimized. The Perfectly Matched Layer of Bérenger (called PML, see [6]) is a convincing solution to this problem in a uniform acoustic medium. Its performance often carries through in a general heterogeneous acoustic medium $c(\mathbf{x})$, though its derivation strictly speaking does not.

PML consists of analytically continuing the solution to the complex plane for points inside the layer. This means we apply a coordinate change of the type $x \to x + i\int^x \sigma(\widetilde{x})\, d\widetilde{x}$, say for a layer in the positive x direction, with $\sigma = 0$ in the interior Ω, and σ positive and increasing inside the layer. Hence the equations are unchanged in the interior, so the solution there is the desired solution to the Helmholtz equation, but waves inside the layer will be more and more damped, the deeper they go into the layer. This is because the solution is a superposition of complex exponentials, which become decaying exponentials under the change of variables. Then we simply put zero Dirichlet boundary conditions at the end of the layer. Whatever waves reflect there will be tiny when they come back out of the layer into the interior. For a heterogeneous

medium, we may still define a layer-based scheme from a transformation of the spatial derivatives which mimics the one done for the PML in a uniform medium, by replacing the Laplacian operator Δ by some Δ_{layer} inside the PML, but this layer will not be perfectly matched anymore and is called a *pseudo-PML* (pPML). In this case, reflections from the interface between Ω and the layer are usually not small. It has been shown in [44] that, in some cases of interest to the optics community with nonuniform media, pPML for Maxwell's equations can still work, but the layer needs to be made very thick in order to minimize reflections at the interface. In this case, the Helmholtz equation has to be solved in a very large computational domain, where most of the work will consist in solving for the pPML. In fact, the layer might even cause the solution to grow exponentially inside it, instead of forcing it to decay ([17], [41]), because the group and phase velocities have an opposite sign.

An ABC scheme which is more stable by construction is the one of Appelö and Colonius [2]. They use a smooth coordinate transform to reduce an unbounded domain to a bounded one with a slowing-down layer, and damp the spurious waves thus created by artificial viscosity (high-order undivided differences). The stability of this scheme follows from its construction, so that it can be used in problems for which the pPML is unstable. However, this method is not ideal because it requires discretizing higher and higher order space derivatives in order to obtain better and better results.

2.2.3 Complexity of ABCs in heterogeneous media

Unfortunately, discrete absorbing layers such as the pPML may need to be quite wide in practice, or may be otherwise computationally costly (because for example of high-order artificial viscosity in [2]). Call L this width (in meters). Although this is not a limitation of the framework presented in this paper, we discretize the Helmholtz operator in the most elementary way using the standard five-point difference stencil. Put $h = 1/N$ for the grid spacing, where N is the number of points per dimension for the interior problem, inside the unit square $\Omega = [0, 1]^2$. While Ω contains N^2 points, the total number of unknowns is $O\left((N + 2w)^2\right)$ in the presence of the layer, where $w = L/h$ is its width in number of grid points. In a uniform medium, the PML width L needed is a fraction of the wavelength, i.e. $L \sim \lambda = \frac{2\pi}{\omega} \sim \frac{1}{N}$, so that we need a constant number of points independently of N: $w = L/h = LN \sim 1$. However, in nonuniform media, the heterogeneity of $c(\mathbf{x})$ can limit the accuracy of the layer. If we consider an otherwise uniform medium with an embedded scatterer outside of Ω, then the pPML will have to be large enough to enclose this scatterer, no matter N. For more general, heterogeneous media such as the ones considered in this paper, we often observe that convergence as a function of L or w is delayed compared to a uniform medium. That means that we have $L \sim L_0$ so that $w \sim NL_0$ or $w = O(N)$, as we assume in the sequel.

The authors of [2] have not investigated the computational complexity of their layer for a given accuracy, but it is clear that a higher accuracy will require higher-order derivatives for the artificial viscosity, and those are quite costly. Fortunately, the framework to be developed over the next chapters also applies to the compression of such a layer, just as it does to any other ABC.

In the case of a second-order discretization, the rate at which one must increase N in order to preserve a constant accuracy in the solution, as ω grows, is about $N = O(\omega^{1.5})$. This unfortunate phenomenon, called the *pollution effect*, is well-known: it begs to *increase* the resolution, or number of points per wavelength, of the scheme as ω grows [4, 3]. As we saw, the width of the pPML may be as wide as a constant value L_0 independent of N, hence its width generally needs to scale as $O(\omega^{1.5})$ grid points.

Next, we introduce the exterior and half-space problems for the Helmholtz equation. We explain how those are related, and how knowledge from the solution of one will help us with the solution to the other.

2.3 The Helmholtz equation: exterior and half-space problems

The previous two sections addressed the fact that we wish to obtain the exterior DtN map in order to approximate the free-space solution in Ω, and how to do that using ABCs and the DtN map. However, ABCs can be computationally intensive. To obtain the exterior DtN map numerically in a feasible way, we will need solve in chapter 3 the exterior problem, and so we define it here. The *half-space problem* for the Helmholtz equation is also interesting to us because we can write down an analytical formula for the DtN map, and use that to gain knowledge that might prove to be more general and apply to the exterior DtN map. Hence we begin by explaining the exterior problem. Then we introduce the half-space problem and its DtN map, and why this might give us insights into the exterior DtN map. We then state important results which will be used in the next chapters.

2.3.1 The exterior problem

The exterior problem consists of solving the free-space problem, but outside of some domain Ω, given a Dirichlet boundary condition g on $\partial\Omega$ and the SRC (2.2). That is, the following has to hold:

$$\Delta u(\mathbf{x}) + \frac{\omega^2}{c^2(\mathbf{x})} u(\mathbf{x}) = 0, \ \mathbf{x} = (x_1, x_2) \in \Omega^c, \tag{2.8}$$

where $\Omega^c = \mathbb{R}^2 \setminus \Omega$, with the boundary data $u(\mathbf{x}) = g(\mathbf{x})$ for $\mathbf{x} \in \partial\Omega$ and a given g. Again, we require the SRC to hold. We call this the *exterior problem*, since we solve for the solution u outside of the domain Ω. We are interested in this problem because, if we can find the solution u, then we can take its derivative, normal to $\partial\Omega$, and obtain the exterior Dirichlet-to-Neumann map. This is how we will calculate the DtN map numerically in the next chapter. Then, of course, the DtN map can be used to solve the free-space problem reformulated on Ω, which is our goal. This sounds like a circular way of solving the free-space problem. This is because, as we shall see in more detail in the next chapter, solving the exterior problem a few times will give

us the exterior DtN map, which will speed up all free-space solves.

In the next chapter, we will use $\Omega = [0,1]^2$, a square of side 1. For practical computations, a rectangular domain made to fit tightly around the scatterer of interest is often used, especially if we have a thin and long scatterer. As we will see in section 2.2, numerous ABCs have been designed for the rectangular domain, so choosing a rectangular domain here is not a limitation. Then, the numerical DtN map is a matrix which, when multiplied by a vector of Dirichlet values u on $\partial\Omega$, outputs Neumann values $\partial_\nu u$ on $\partial\Omega$. In particular, we can consider the submatrix of the DtN map corresponding to Dirichlet values on one particular side of Ω, and Neumann values on that same side. As we shall see next, this submatrix should be quite similar to the half-space DtN map.

2.3.2 The half-space problem

We consider again the scalar Helmholtz equation, but this time we take Ω to be the top half-space $\mathbb{R}^2_+ = \{(x_1, x_2) : x_2 \geq 0\}$, so that now the boundary $\partial\Omega$ is the x_1-axis, that is, when $x_2 = 0$. And we consider the exterior problem for this Ω:

$$\Delta u(\mathbf{x}) + \frac{\omega^2}{c^2(\mathbf{x})} u(\mathbf{x}) = 0, \qquad \mathbf{x} = (x_1, x_2), \qquad x_2 \leq 0 \tag{2.9}$$

with given boundary condition g:

$$u(x_1, 0) = g(x_1, 0), \tag{2.10}$$

requiring some decay on $g(x_1, 0)$ as $x_1 \to \pm\infty$ and the SRC (2.2) to hold in the bottom half-space $\mathbb{R}^2_- = \{(x_1, x_2) : x_2 \leq 0\}$. We shall explain how to find this analytical DtN map for the half-space problem in uniform medium, but first, we discuss the relevance of the half-space DtN map for the exterior DtN map.

Let us use again $\Omega = [0,1]^2$ as we introduced earlier, and let us call S_1 the bottom side of $\partial\Omega$: $S_1 = \{(x_1, x_2) : 0 \leq x_1 \leq 1, x_2 = 0\}$. For the exterior problem, we prescribe boundary values g on $\partial\Omega$, solve the exterior problem and obtain values of u_{ext} everywhere outside of Ω. From those, we know the exterior DtN map $D_{\text{ext}} : \partial\Omega \to \partial\Omega$. Consider now the values of u_{ext} we have just found, along the x_1 axis: we can use those to define the boundary condition (2.10) of the half-space problem. The solution u_{half} we find for this half-space problem on the bottom half-space $\mathbb{R}^2_-$ coincides with u_{ext} on that half-space. Similarly, the exterior DtN map D_{ext} restricted to the bottom side of Ω, $D_{\text{ext}} : S_1 \to S_1$ coincides with the half-space DtN map D_{half}, restricted to this same side $D_{\text{half}} : S_1 \to S_1$.

This relationship between the half-space and exterior DtN maps remains when solving those problems numerically. To solve the exterior problem numerically, we proceed very similarly to how we would for the free-space problem. As we saw in the previous section on ABCs, we may place an ABC on $\partial\Omega$ for the free-space problem. For the exterior problem, we place this ABC just a little outside of $\partial\Omega$, as in Figure 2-1. We then enforce the boundary condition $u(\mathbf{x}) = g(\mathbf{x})$ for $\mathbf{x}$ on $\partial\Omega$, and solve the

Helmholtz equation outside of Ω, inside the domain delimited by the ABC. We thus obtain the solution u on a thin strip just outside of Ω, and we can compute from that u the DtN map $\partial_\nu u$ on $\partial\Omega$. This is why we need to put the ABC just a little outside of Ω, and not exactly on $\partial\Omega$.

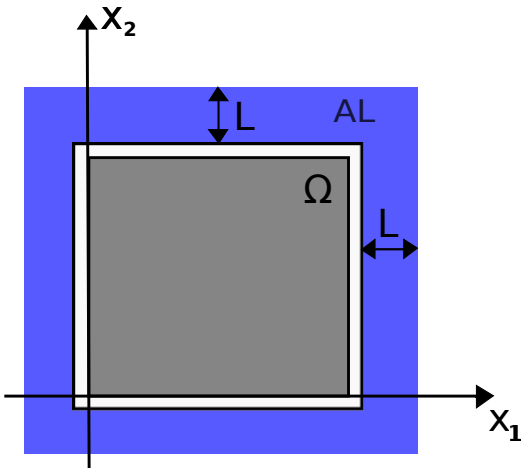

Figure 2-1: The exterior problem: Ω is in grey, there is a thin white strip around Ω, then an absorbing layer in blue.

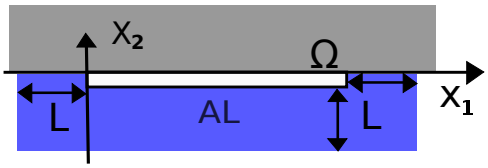

Figure 2-2: The half-space problem: there is thin white strip below the bottom side of Ω, and an absorbing layer (in blue) surrounds that strip on three sides.

To solve the half-space problem numerically, we will need again an ABC, in order to reformulate the problem on a smaller, bounded domain. We can prescribe values of u only along the bottom edge S_1 as in Figure 2-2, leave a thin strip just below that edge, surround that thin strip by an ABC on all three other sides, and solve for the solution u. We can then, from the solution u in this thin strip, calculate an approximate half-space DtN map on S_1. And we see how this approximate half-space DtN map restricted to S_1 will be similar (but not exactly the same) to the exterior DtN map restricted to that same edge S_1, given the same boundary data on S_1 and 0 boundary data on the other edges of $\partial\Omega$. The main difference between the two maps is that in the exterior case, the two corners of S_1 will cause scattering, some of which will affect the solution in the top half-space $\mathbb{R}^2_+$ as well.

It is because of this connection between the half-space and the exterior DtN maps that we analyze further the half-space DtN map, and use insights obtained from this analysis to treat the exterior DtN map.

2.3.3 Solution to the half-space problem in a uniform medium using the Green's function

In section 2.1.1 we used the free-space Green's function G to obtain the solution u anywhere from integrating G with the right-hand side f. We can do the same for the

half-space problem. First, we define the Green's function of the half-space problem:

Definition 3. *The half-space* Green's function *for the Helmholtz equation 2.9 is the unique function $G(\mathbf{x}, \mathbf{y})$ which solves the half-space problem with zero boundary data, that is $g = 0$ in (2.10), and with right-hand side the delta function $f(\mathbf{x}) = \delta(|\mathbf{x} - \mathbf{y}|)$, for every fixed $\mathbf{y}$.*

This half-space Green's function, which we shall call G_{half}, is once again well-known, and can be obtained from the free-space Green's function by the reflection principle (p. 110 of [24]), with $\mathbf{x} = (x_1, x_2)$, $\mathbf{y} = (y_1, y_2)$ and $\mathbf{x}' = (-x_1, x_2)$:

$$G_{\text{half}}(\mathbf{x}, \mathbf{y}) = G(\mathbf{x}, \mathbf{y}) - G(\mathbf{x}', \mathbf{y}). \tag{2.11}$$

Then, the solution u to the half-space problem with $g = 0$ is as expected:

$$u(\mathbf{x}) = \int_S G_{\text{half}}(|\mathbf{x} - \mathbf{y}|) f(\mathbf{y})\ dy,\ \mathbf{x} \in \mathbb{R}^2 \setminus S. \tag{2.12}$$

where $S = \{(x_1, x_2) : x_2 = 0\}$ is the x_1 axis. This half-space Green's function will be useful to us, and in particular, we are interested in the following result of [20], slightly reformulated for our purposes:

Lemma 1. Theorem 2.3 of [20]. *Let G_{half} be the half-space Green's function as defined above. Let $n \in \mathbb{N}$ be some discretization parameter, n even, and let $h = 1/n$. Let $Y = \{\mathbf{y}_j = (jh, -h), j = 1, \ldots, n/2\}$ and $X = \{\mathbf{x}_j = (jh, -h), j = n/2 + 1, \ldots, n\}$. Then $\left(G_{half}(\mathbf{x}, \mathbf{y})\right)_{\mathbf{x} \in X,\ \mathbf{y} \in Y}$ is numerically low-rank. More precisely, for any $\varepsilon > 0$, there exist a constant $J = O(\log k |\log \varepsilon|^2)$, functions $\{\alpha_p(\mathbf{x})\}_{1 \le p \le J}$ for $\mathbf{x} \in X$ and functions $\{\beta_p(\mathbf{y})\}_{1 \le p \le J}$ for $\mathbf{y} \in Y$ such that*

$$\left| G(\mathbf{x}, \mathbf{y}) - \sum_{p=1}^{J} \alpha_p(\mathbf{x}) \beta_p(\mathbf{y}) \right| \le \varepsilon \text{ for } \mathbf{x} \in X, \mathbf{y} \in Y.$$

We provide this lemma here because the concept of an operator being *numerically low-rank* away from its diagonal will be useful later.

For now, we define the half-space DtN map. We first want to find a kernel that will give us the solution on the bottom half-space when multiplied with the boundary data g and not with the right-hand side f, which we will soon assume to be 0. Such a kernel exists, and is related to G_{half}. From equation (2.37) of [24], we may write

$$u(\mathbf{x}) = \int_S \partial_{\nu_{\mathbf{y}}} G_{\text{half}}(\mathbf{x}, \mathbf{y})\ u(\mathbf{y})\ dS_{\mathbf{y}},\ \mathbf{x} \in \mathbb{R}^2 \setminus S$$

where ν is the outward- (thus downward-) pointing normal to S. Hence we obtain

$$u(\mathbf{x}) = -\int \partial_{y_2} G_{\text{half}}\left(\mathbf{x}, (y_1, y_2)\right)\big|_{y_2 = 0}\ u(y_1, 0)\ dy_1,\ \mathbf{x} \in \mathbb{R}^2 \setminus S. \tag{2.13}$$

We have found a different way of expressing the solution u to the half-space problem using the half-space Green's function G_{half}.

2.3.4 The analytical half-space DtN map

Since we wish to consider the exterior Dirichlet-to-Neumann map for u on $x_2 = 0$, $Du(x_1, 0) = \partial_{\nu_{\mathbf{x}}} u(x_1, x_2)|_{x_2=0} = -\partial_{x_2} u(x_1, x_2)|_{x_2=0}$, we take the normal derivative in (2.13) and evaluate at $x_2 = 0$ (using the fact that $u(y_1, 0) = g(y_1, 0)$ is the boundary data):

$$Du(x_1, 0) = \int \partial_{x_2}\partial_{y_2} G_{\text{half}}\left((x_1, x_2), (y_1, y_2)\right)|_{x_2=0, y_2=0}\ g(y_1, 0)\ dy_1,\ x_1 \in \mathbb{R}.$$

Hence we have found the kernel of the half-space DtN map to be, essentially, two derivatives of the half-space Green's function. Since we know that Green's function, we can use it to obtain an analytical expression for the half-space DtN map. Notice that $\frac{\partial}{\partial z} H_0^{(1)}(z) = -H_1^{(1)}(z)$, We have

$$\partial_{x_2}\partial_{y_2} G(\mathbf{x}, \mathbf{y}) = \frac{i}{4}\partial_{x_2}\partial_{y_2} H_0^{(1)}(z)$$

with

$$z = k|\mathbf{x} - \mathbf{y}| = k\sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2},$$

so that

$$\partial_{y_2} z = k^2 \frac{y_2 - x_2}{z},\ \partial_{x_2} z = k^2 \frac{x_2 - y_2}{z}.$$

Hence

$$\partial_{y_2} G(\mathbf{x}, \mathbf{y}) = \frac{i}{4}\frac{\partial z}{\partial y_2}\partial_z H_0^{(1)}(z) = \frac{ik^2}{4}\frac{x_2 - y_2}{z} H_1^{(1)}(z)$$

and

$$\partial_{x_2}\partial_{y_2} G(\mathbf{x}, \mathbf{y}) = \frac{ik^2}{4}\left(\frac{z - k^2(x_2 - y_2)^2/z}{z^2} H_1^{(1)}(z) + k^2\left(\frac{x_2 - y_2}{z}\right)^2 \partial_z H_1^{(1)}(z)\right). \tag{2.14}$$

Also, we have

$$\partial_{x_2}\partial_{y_2} G(\mathbf{x}', \mathbf{y}) = \frac{i}{4}\partial_{x_2}\partial_{y_2} H_0^{(1)}(z')$$

with

$$z' = k\sqrt{(x_1 - y_1)^2 + (-x_2 - y_2)^2},$$

so that

$$\partial_{y_2} z' = k^2 \frac{y_2 + x_2}{z'},\ \partial_{x_2} z' = k^2 \frac{y_2 + x_2}{z'}.$$

Hence

$$\partial_{y_2} G(\mathbf{x}', \mathbf{y}) = \frac{i}{4}\frac{\partial z'}{\partial y_2}\partial_{z'} H_0^{(1)}(z') = \frac{ik^2}{4}\frac{-x_2 - y_2}{z'} H_1^{(1)}(z')$$

and

$$\partial_{x_2}\partial_{y_2}G(\mathbf{x}',\mathbf{y}) = \frac{ik^2}{4}\left(\frac{-z' + k^2(x_2+y_2)^2/z'}{z'^2}H_1^{(1)}(z') - k^2\left(\frac{x_2+y_2}{z'}\right)^2 \partial_{z'}H_1^{(1)}(z')\right). \tag{2.15}$$

Now let $x_2 = 0, y_2 = 0$ but $0 \neq k|\mathbf{x}-\mathbf{y}| = k|x_1 - y_1| = z|_{x_2=0,y_2=0} = z'|_{x_2=0,y_2=0}$ in (2.14) and (2.15), so that

$$\begin{aligned}
\partial_{x_2}\partial_{y_2}\, G_{\text{half}}(\mathbf{x},\mathbf{y})|_{x_2=0,y_2=0} &= \partial_{x_2}\partial_{y_2}\, G(\mathbf{x},\mathbf{y})|_{x_2=0,y_2=0} - \partial_{x_2}\partial_{y_2}\, G(\mathbf{x}',\mathbf{y})|_{x_2=0,y_2=0} \\
&= \frac{ik^2}{2k|\mathbf{x}-\mathbf{y}|}H_1^{(1)}(k|\mathbf{x}-\mathbf{y}|)\Big|_{x_2=0,y_2=0} \\
&= \frac{ik^2}{2k|x_1-y_1|}H_1^{(1)}(k|x_1-y_1|)
\end{aligned}$$

Thus we have found that the half-space DtN map kernel is:

$$K(r) = \frac{ik^2}{2kr}H_1^{(1)}(kr) \tag{2.16}$$

and thus the half-space DtN map is:

$$Du(x_1,0) = \int K(|x_1-y_1|)g(y_1,0)\, dy_1. \tag{2.17}$$

As we shall prove in chapter 4, the half-space DtN map kernel (2.16) for a uniform medium is numerically low-rank away from its singularity, just as the half-space Green's function is from Lemma 1. This means that, if x and y are coordinates along the infinite boundary S, and $|x-y| \geq r_0$ for some constant quantity r_0, then the DtN map, a function of x and y, can be approximated up to error ε as a sum of products of functions of x only with functions of y only (*separability*) and that this sum is finite and in fact involves only a small numbers of summands (*low-rank*). Hence, a numerical realization of that DtN map in matrix form should be compressible. In particular, blocks of that matrix which are away from the diagonal should have low column rank. We shall make all of this precise in chapter 4.

Recall that the goal of this thesis, to compress ABCs, will be achieved by approximating the DtN map in more general cases than the uniform half-space problem. Before we explain our approach for compressing an ABC in the next two chapters, we first explain the most straightforward way of obtaining a DtN map from an ABC, by eliminating the unknowns in the absorbing layer in order to obtain a reduced system on the interior nodes. This solution, however, is computationally impractical. It only serves to make explicit the relationship between ABCs and the DtN map.

2.4 Eliminating the unknowns: from any ABC to the DtN map

We now wish to explain the fundamental relationship between discrete ABCs and the discrete DtN map: any discrete ABC can be reformulated as a discrete DtN map on $\partial\Omega$. We present two similar ways of obtaining that relationship using Schur complements: one for the half-space problem, and one for the free-space problem.

2.4.1 Eliminating the unknowns in the half-space problem

We consider the half-space problem, in which we care about the solution in the top half-plane. We assume from the SRC that the solution far away from $x_2 = 0$ in the bottom half is small. We want to eliminate unknowns from far away on the bottom side, where the solution is so small we ignore it as zero, towards the positive x_2 direction in order to obtain an outgoing Dirichlet-to-Neumann map in the $-x_2$ direction. We assume $f = 0$ everywhere in the bottom half-plane. Let u_1 denote the first line of unknowns at the far bottom, u_2 the next line, and so on. We first use a Schur complement to eliminate u_1 from the discretized (using the standard five-point stencil) Helmholtz system which is as follows:

$$\frac{1}{h^2}\begin{pmatrix} S_1 & P_1 \\ P_1^T & C_1 \end{pmatrix} \begin{pmatrix} u_1 \\ \vdots \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \end{pmatrix}.$$

We then define similarly the matrices S_k, P_k and C_k corresponding to having eliminated lines u_1 through u_{k-1} from the system. Only the upper block of C_1, or C_k when subsequently eliminating line u_k, will be modified from the Schur complement. Indeed, since we eliminate from bottom to top and we have a five-point stencil, the matrices P_k will be

$$P_k = \begin{pmatrix} I & 0 & \cdots \end{pmatrix}$$

and so we may write the following recursion for the Schur complements:

$$S_{k+1} = M - S_k^{-1}, \tag{2.18}$$

where M is the main block of the 2D Helmholtz operator multiplied by h^2, that is,

$$M = \begin{pmatrix} -4+h^2k^2 & 1 & \\ 1 & -4+h^2k^2 & \ddots \\ & \ddots & \ddots \\ & & \end{pmatrix} = -2I + Lh^2.$$

Here L is the 2nd order discretization of the 1D Helmholtz operator in the x_1 direction, $\partial_{x_1}\partial_{x_1} + k^2$. Now, at each step we denote

$$S_k = hD_k - I. \tag{2.19}$$

Indeed, looking at the first block row of the reduced system at step $k-1$ we have

$$S_k u_k + I u_{k+1} = 0$$

or

$$(hD_k - I)u_k + I u_{k+1} = 0.$$

From this we obtain the DtN map D_k from a backward difference:

$$\frac{u_k - u_{k+1}}{h} = D_k u_k.$$

Now we use (2.19) inside (2.18) to obtain

$$hD_{k+1} - I = M + (I - hD_k)^{-1} = M + I + hD_k + h^2 D_k^2 + O(h^3)$$

or

$$hD_{k+1} - hD_k = Lh^2 + h^2 D_k^2 + O(h^3),$$

which we may rewrite to obtain a discretized Riccati equation (something similar was done in [36]):

$$\frac{D_{k+1} - D_k}{h} = L + D_k^2 + O(h), \tag{2.20}$$

of which we may take the limit as $h \to 0$ to obtain the Riccati equation for the DtN map D in the $-x_2$ direction:

$$D_{x_2} = (\partial_{x_1}\partial_{x_1} + k^2) + D^2. \tag{2.21}$$

This equation is to be evolved in the $+x_2$ direction, starting far away in the bottom half-space. Looking at the steady state, $D_{x_2} = 0$, we get back $D^2 = -\partial_{x_1}\partial_{x_1} - k^2$, which is the Helmholtz equation with 0 right-hand side f (which we have assumed to hold in the bottom half-space). Hence we conclude that the Riccati equation for the DtN map could be used to obtain the DtN map in the half-space case, and maybe even more complicated problems. We leave this to future work, and turn to a very similar way of eliminating unknowns, but for the exterior DtN map with $\Omega = [0,1]^2$ this time. This technique will not give rise to a Riccati equation, but will help us understand how the DtN map can be used numerically to solve the free-space Helmholtz equation reformulated on Ω.

2.4.2 Eliminating the unknowns in the exterior problem

In this subsection, we assume we need an absorbing layer of large width, $w \geq N$ in number of points. We write the system for the discrete Helmholtz equation as

$$\begin{pmatrix} A & P \\ P^T & C \end{pmatrix} \begin{pmatrix} u_{out} \\ u_\Omega \end{pmatrix} = \begin{pmatrix} 0 \\ f_\Omega \end{pmatrix}, \tag{2.22}$$

with $A = \Delta_{layer} + k^2 I$ and $C = \Delta + k^2 I$, with Δ overloaded to denote discretization of the Laplacian operator, and Δ_{layer} the discretization of the Laplacian operator inside the absorbing layer. We wish to eliminate the exterior unknowns u_{out} from this system in order to have a new system which only depends on the interior unknowns u_Ω. The most obvious way of eliminating those unknowns is to form the Schur complement $S = C - P^T A^{-1} P$ of A by any kind of Gaussian elimination. For instance, in the standard raster scan ordering of the unknowns, the computational cost of this method[4] is $O(w^4)$ — owing from the fact that A is a sparse banded matrix of size $4Nw + 4w^2$ which is $O(w^2)$, and band $N + 2w$. Alternatively, elimination of the unknowns can be performed by layer-stripping, starting with the outermost unknowns from u_{out}, until we eliminate the layer of points that is just outside of $\partial\Omega$. The computational cost will be $O(w^4)$ in this case as well. To see this, let u_w be the points on the outermost layer, u_{w-1} the points in the layer just inside of u_w, etc. Then we have the following system:

$$\begin{pmatrix} A_w & P_w \\ P_w^T & C_w \end{pmatrix} \begin{pmatrix} u_w \\ \vdots \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \end{pmatrix}$$

Note that, because of the five-point stencil, P_w has non-zeros exactly on the columns corresponding to u_{w-1}. Hence the matrix $P_w^T A_w^{-1} P_w$ in the first Schur complement $S_w = C_w - P_w^T A_w^{-1} P_w$ is non-zero exactly at the entries corresponding to u_{w-1}. It is then clear that, in the next Schur complement, to eliminate the next layer of points, the matrix A_{w-1} (the block of S_w corresponding to the points u_{w-1}) to be inverted will be full. For the same reason, every matrix A_j to be inverted thereafter, for every subsequent layer to be eliminated, will be a full matrix. Hence at every step the cost of forming the corresponding Schur complement is at least on the order of m^3, where m is the number of points in that layer. Hence the total cost of eliminating

[4]The cost of the Schur complement procedure is dominated by that of Gaussian elimination to apply A^{-1} to P. Gaussian elimination on a sparse banded matrix of size s and band b is $O(sb^2)$, as can easily be infered from Algorithm 20.1 of [51].

the exterior unknowns by layer stripping is approximately

$$\sum_{j=1}^{w}(4(N+2j))^3 = O(w^4).$$

Similar arguments can be used for the Helmholtz equation in 3 dimensions. In this case, the computational complexity of the Schur complement or layer-stripping methods would be $O(w^3(w^2)^2) = O(w^7)$ or $\sum_{j=1}^{w}(6(N+2j)^2)^3 = O(w^7)$. Therefore, direct elimination of the exterior unknowns is quite costly. Some new insight will be needed to construct the DtN map more efficiently.

We now remark that, whether we eliminate exterior unknowns in one pass or by layer-stripping, we obtain a reduced system. It looks just like the original Helmholtz system on the interior unknowns u_Ω, except for the top left block, corresponding to u_0 the unknowns on $\partial\Omega$, which has been modified by the elimination procedure. Hence with the help of some dense matrix D we may write the reduced, N^2 by N^2 system as

$$Lu = \begin{pmatrix} (hD-I)/h^2 & I/h^2 & 0 & \cdots \\ I/h^2 & & & \\ 0 & & [\,\Delta + k^2 I\,] & \\ \vdots & & & \end{pmatrix} \begin{pmatrix} u_0 \\ u_{-1} \\ u_{-2} \\ \vdots \end{pmatrix} = \begin{pmatrix} 0 \\ f_{-1} \\ f_{-2} \\ \vdots \end{pmatrix} \tag{2.23}$$

and we have thus obtained an absorbing boundary condition which we may use on the boundary of Ω, independent of the right-hand side f. Indeed, if we call u_{-1} the first layer of points inside Ω, we have $(I - hD)u_0 = u_{-1}$, or

$$\frac{u_0 - u_{-1}}{h} = Du_0,$$

a numerical realization of the DtN map in (2.6), using the ABC of choice. Indeed, elimination can be used to reformulate any computationally intensive ABC, not just absorbing layers, into a realization of (2.6). Any ABC is equivalent to a set of equations relating unknowns on the surface to unknowns close to the surface, and possibly auxiliary variables. Again, elimination can reduce those equations to relations involving only unknowns on the boundary and on the first layer inside the boundary, to obtain a numerical DtN map D. A drawback is that forming this matrix D by elimination is prohibitive, as we have just seen.

As for the complexity of solving the Helmholtz equation, reducing the ABC confers the advantage of making the number of nonzeros in the matrix L (of Section 2.4) independent of the width of the absorbing layer or complexity of the ABC. After elimination of the layer, it is easy to see that L has about $20N^2$ nonzero entries, instead

of the $5N^2$ one would expect from a five-point stencil discretization of the Helmholtz equation, because the matrix D (part of a small block of L) is full. Although obtaining a fast matrix-vector product for our approximation of D could reduce the application cost of L from $20N^2$ to something closer to $5N^2$, it should be noted that the asymptotic complexity does not change – only the constant does.

This thesis addresses those two problems, obtaining the DtN map and applying it fast. The next chapter, chapter 3, suggests adapting the framework of matrix probing in order to obtain D in reasonable complexity. Subsequently, chapter 4 presents a compression method which leads to a fast application of the DtN map.

Chapter 3

Matrix probing for expanding the Dirichlet-to-Neumann map

Recall that the goal of this thesis is to introduce a new compression scheme for ABCs. This scheme consists of two steps:

1. a precomputation sets up an expansion of the Dirichlet-to-Neumann map, then
2. a fast algorithm is used to apply the DtN map in a Helmholtz solver.

This chapter is concerned with the first step of this procedure, namely, setting up an expansion of the exterior DtN map in a precomputation. This will pave the way for compression in step two, presented in the next chapter.

The main strategy we use in this chapter is matrix probing, we introduce it in section 3.1. For matrix probing to be an efficient expansion scheme, we need to carefully choose the basis for this expansion. We present our choices and their rationales in section 3.2. In particular, inverse powers multiplied by a complex exponential work quite well as kernels for the basis. We then present a detailed study of using matrix probing to expand the DtN map in various different media, use that expansion to solve the Helmholtz equation, and document the complexity of the method, all in section 3.3. Then, we prove in section 3.4 a result on approximating the half-space DtN map with the particular set of functions mentioned before, inverse powers multiplied by a complex exponential. We also present a numerical confirmation of that result in section 3.5.

3.1 Introduction to matrix probing

The idea of matrix probing is that a matrix D with adequate structure can sometimes be recovered from the knowledge of a fixed, small number of matrix-vector products Dg_j, where g_j are typically random vectors. In the case where D is the numerical DtN map (with a slight abuse of notation), each g_j consists of Dirichlet data on $\partial\Omega$, and each application Dg_j requires solving an exterior Helmholtz problem to compute the derivative of the solution normal to $\partial\Omega$. We first explain how to obtain the

matrix-vector multiplication of the DtN map with any vector, without having to use the costly procedure of layer-stripping. We then introduce matrix probing.

3.1.1 Setup for the exterior problem

Recall the exterior problem of section 2.3.1: solving the heterogeneous-medium Helmholtz equation at frequency ω, outside $\Omega = [0,1]^2$, with Dirichlet boundary condition $u = g$ on $\partial\Omega$. This problem is solved numerically with the five-point stencil of finite differences (FD), using h to denote the grid spacing and N the number of points across one dimension of Ω. We use a Perfectly Matched Layer (PML) or pPML, introduced in section 2.2.2, as our ABC. The layer starts at a fixed, small distance away from Ω, so that we keep a small strip around Ω where the equations are unchanged. Recall that the width of the layer is in general as large as $O(\omega^{1.5})$ grid points. We number the edges of $\partial\Omega$ counter-clockwise starting from $(0,0)$, hence side 1 is the bottom edge $(x,0)$, $0 \le x \le 1$, side 2 is the right edge, etc. The exterior DtN map for this problem is defined from $\partial\Omega$ to itself. Thus its numerical realization, which we also call D by a slight abuse of notation, has a 4×4 block structure. Hence the numerical DtN map D has 16 sub-blocks, and is $n \times n$ where $n = 4N$. As an integral kernel, D would have singularities at the junctions between these blocks (due to the singularities in $\partial\Omega$), so we shall respect this feature by probing D sub-block by sub-block. We shall denote a generic such sub-block by M, or as the (i_M, j_M) sub-block of D, referring to its indices in the 4×4 sub-block structure.

The method by which the system for the exterior problem is solved is immaterial in the scope of this paper, though for reference, the experiments in this paper use UMFPACK's sparse direct solver [14]. For treating large problems, a better solver should be used, such as the sweeping preconditioner of Engquist and Ying [20, 19], the shifted Laplacian preconditioner of Erlangga [22], the domain decomposition method of Stolk [49], or the direct solver with spectral collocation of Martinsson, Gillman and Barnett [26, 25]. This in itself is a subject of ongoing research which we shall not discuss further.

For a given boundary condition g, we solve the system and obtain a solution u in the exterior computational domain. In particular we consider u_1, the solution in the layer just outside of $\partial\Omega$. We are using the same notation as in section 2.4, where as we recall u_0 was the solution on the boundary, hence here $u_0 = g$. We know from Section 2.4 that u_1 and g are related by

$$\frac{u_1 - g}{h} = Dg \tag{3.1}$$

The matrix D that this relation defines needs not be interpreted as a first-order approximation of the continuous DtN map: it is the algebraic object of interest that will be "probed" from repeated applications to different vectors g.

Similarly, for probing the (i_M, j_M) block M of D, one needs matrix-vector products of D with vectors g of the form $[z,0,0,0]^T$, $[0,z,0,0]^T$, etc., to indicate that the Dirichlet boundary condition is z on the side indexed by j_M, and zero on the other sides. The application Dg is then restricted to side i_M.

3.1.2 Matrix probing

The dimensionality of D needs to be limited for recovery from a few Dg_j to be possible, but matrix probing is *not* an all-purpose low-rank approximation technique. Instead, it is the property that D has an efficient representation in some adequate pre-set basis that makes recovery from probing possible. As opposed to the randomized SVD method which requires the number of matrix-vector applications to be greater than the rank [32], matrix probing can recover interesting structured operators from a single matrix-vector application [12, 15].

We now describe a model for M, any $N \times N$ block of D, that will sufficiently lower its dimensionality to make probing possible. Assume we can write M as

$$M \approx \sum_{j=1}^{p} c_j B_j \tag{3.2}$$

where the B_j's are fixed, known basis matrices, that need to be chosen carefully in order to give an accurate approximation of M. In the case when the medium c is uniform, we typically let B_j be a discretization of the integral kernel

$$B_j(x, y) = \frac{e^{ik|x-y|}}{(h + |x - y|)^{j/2}}, \tag{3.3}$$

where again $h = 1/N$ is the discretization parameter. We usually add another index to the B_j, and a corresponding multiplicative factor, to allow for a smooth dependence on $x + y$ as well. We shall further detail our choices and discuss their rationales in Section 3.2. For now, we note that the advantage of the specific choice of basis matrix (3.3), and its generalizations explained in Section 3.2, is that it results in accurate expansions with a number of parameters p which is "essentially independent" of N, namely that grows either logarithmically in N, or at most like a very sublinear fractional power law (such as $N^{0.12}$, see section 3.3.4). This is in sharp contrast to the scaling for the layer width, $w = O(N)$ grid points, discussed earlier. The form of B_j suggested in equation (3.3) is motivated by the fact that they provide a good expansion basis for the uniform-medium half-space DtN map in $\mathbb{R}^2$. This will be proved in section 3.4.

Given a random vector $z^{(1)} \sim N(0, I_N)$ (other choices are possible), the product $w^{(1)} = Mz^{(1)}$ and the expansion (3.2), we can now write

$$w^{(1)} = Mz^{(1)} \approx \sum_{j=1}^{p} c_j B_j z^{(1)} = \Psi_{z^{(1)}} \mathbf{c}. \tag{3.4}$$

Multiplying this equation on the left by the pseudo-inverse of the N by p matrix $\Psi_{z^{(1)}}$ will give an approximation to $\mathbf{c}$, the coefficient vector for the expansion (3.2) of M. More generally, if several applications $w^{(j)} = Mz^{(j)}$, $j = 1, \ldots, q$ are available, a larger system is formed by concatenating the $\Psi_{z^{(j)}}$ into a tall-and-thin Nq by p matrix $\mathbf{\Psi}$. The computational work is dominated, here and in other cases [12, 15], by

the matrix-vector products $Dg^{(1)}$, or $Mz^{(j)}$. Note that both $\Psi_{z^{(j)}}$ and the resulting coefficient vector $\mathbf{c}$ depend on the vectors $z^{(j)}$. In the sequel we let $z^{(j)}$ be gaussian iid random.

In a nutshell, recovery of $\mathbf{c}$ works under mild assumptions on B_j, and when p is a small fraction of Nq up to log factors. In order to improve the conditioning in taking the pseudo-inverse of the matrix Ψ_z and reduce the error in the coefficient vector $\mathbf{c}$, one may use $q > 1$ random realizations of M. There is a limit to the range of p for which this system is well-posed: past work by Chiu and Demanet [12] covers the precise conditions on p, N, and the following two parameters, called *weak condition numbers*, for which recoverability of $\mathbf{c}$ is accurate with high probability.

Definition 4. Weak condition number λ.

$$\lambda = \max_j \frac{\|B_j\|_2 \sqrt{N}}{\|B_j\|_F}$$

Definition 5. Weak condition number κ.

$$\kappa = cond(\mathbf{B}),\ \mathbf{B}_{j\ell} = Tr(B_j^T B_\ell)$$

It is desirable to have a small λ, which translates into a high rank condition on the basis matrices, and a small κ, which translates into a Riesz basis condition on the basis matrices. Having small weak condition numbers will guarantee a small failure probability of matrix probing and a bound on the condition number of $\mathbf{\Psi}$, i.e. guaranteed accuracy in solving for $\mathbf{c}$. Also, using $q > 1$ allows to use a larger p, to achieve greater accuracy. These results are contained in the following theorem.

Theorem 1. *(Chiu-Demanet, [12]) Let z be a Gaussian i.i.d. random vector of length qN, and $\mathbf{\Psi}$ as above. Then $cond(\mathbf{\Psi}) \leq 2\kappa + 1$ with high probability provided that p is not too large, namely*

$$qN \geq C\, p\, (\kappa\lambda \log N)^2,$$

for some number $C > 0$.

As noted previously, the work necessary for probing the matrix M is on the order of q solves of the original problem. Indeed, computing $Mz^{(1)}, \ldots, Mz^{(q)}$ means solving q times the exterior problem with the AL. This is roughly equivalent to solving the original Helmholtz problem with the AL q times, assuming the AL width w is at least as large as N. Then, computing the qp products of the p basis matrices with the q random vectors amounts to a total of at most qpN^2 work, or less if the basis matrices have a fast matrix-vector product. And finally, computing the pseudo-inverse of $\mathbf{\Psi}$ has cost Nqp^2. Hence, as long as $p, q \ll N$, the dominant cost of matrix probing[1] comes from solving q times the exterior problem with a random Dirichlet boundary

[1]We will see later that we also need to perform a QR factorization on the basis matrices, and this has cost N^2p^2. This precomputation has a cost similar or smaller to the cost of an exterior solve using current Helmholtz solvers. It might also be possible to not need a QR factorization if basis matrices closer to orthonormal are used.

condition. In our experiments, $q = O(1)$ and p can be as large as a few hundreds for high accuracy.

Finally, we note that the information from the q solves can be re-used for any other block which is in the same block column as M. However, if it is needed to probe blocks of D which are not all in the same block column, then another q solves need to be performed, with a Dirichlet boundary condition on the appropriate side of $\partial\Omega$. This of course increases the total number of solves. Another option would be to probe all of D at once, using a combination of basis matrices that have the same size as D, but that are 0 except on the support of each distinct block in turn. In this case, κ remains the same because we still orthogonalize our basis matrices, but λ doubles ($\|B_j\|_2$ and $\|B_j\|_F$ do not change but $N \to 4N$) and this makes the conditioning worse, in particular a higher value of q is needed for the same accuracy, given by p. Hence we have decided not to investigate further this approach, which might become more advantageous in the case of a more complicated polygonal domain.

3.1.3 Solving the Helmholtz equation with a compressed ABC

Once we have obtained approximations $\widetilde{M}$ of each block M in compressed form through the coefficients $\mathbf{c}$ using matrix probing, we construct block by block the approximation $\widetilde{D}$ of D and use it in a solver for the Helmholtz equation on the domain $\Omega = [0,1]^2$, with the boundary condition

$$\frac{\partial u}{\partial \nu} = \widetilde{D}u, \qquad x \in \partial\Omega.$$

3.2 Choice of basis matrices for matrix probing

The essential information of the DtN map needs to be summarized in broad strokes in the basis matrices B_j, with the details of the numerical fit left to the probing procedure. In the case of D, most of its physics is contained in its *diagonal singularity* and *oscillations*, as predicted by geometrical optics.

A heuristic argument to obtain the form of D starts from the Green's formula (2.5), that we differentiate one more time in the normal direction. After accounting for the correct jump condition, we get an alternative Steklov-Poincare identity, namely

$$D = (T^* + \frac{1}{2}I)^{-1}H,$$

where H is the hypersingular integral operator with kernel $\frac{\partial^2 G}{\partial \nu_{\mathbf{x}} \partial \nu_{\mathbf{y}}}$, again $G(\mathbf{x}, \mathbf{y})$ is the free-space Green's function and $\nu_{\mathbf{x}}$, $\nu_{\mathbf{y}}$ are the normals to $\partial\Omega$ in $\mathbf{x}$ and $\mathbf{y}$ respectively. The presence of $(T^* + \frac{1}{2}I)^{-1}$ is somewhat inconsequential to the form of D, as it involves solving a well-posed second-kind integral equation. As a result, the properties of D are qualitatively similar to those of H. (The exact construction of D from G is of course already known in a few special cases, such as the uniform medium half-space problem considered earlier.)

3.2.1 Oscillations and traveltimes for the DtN map

Geometrical optics will reveal the form of G. In a context where there is no multi-pathing, that is, where there is a single traveltime $\tau(\mathbf{x}, \mathbf{y})$ between any two points $\mathbf{x}, \mathbf{y} \in \Omega$, one may write a high-$\omega$ asymptotic series for G as

$$G(\mathbf{x}, \mathbf{y}) \sim e^{i\omega\tau(\mathbf{x},\mathbf{y})} \sum_{j\geq 0} A_j(\mathbf{x}, \mathbf{y})\omega^{-j}, \tag{3.5}$$

$\tau(\mathbf{x}, \mathbf{y})$ is the traveltime between points $\mathbf{x}$ and $\mathbf{y}$, found by solving the Eikonal equation

$$\|\nabla_{\mathbf{x}}\tau(\mathbf{x}, \mathbf{y})\| = \frac{1}{c(\mathbf{x})}, \tag{3.6}$$

and the amplitudes A_j satisfy transport equations. In the case of multi-pathing (possible multiple traveltimes between any two points), the representation (3.5) of G becomes instead

$$G(\mathbf{x}, \mathbf{y}) \sim \sum_j e^{i\omega\tau_j(\mathbf{x},\mathbf{y})} \sum_{k\geq 0} A_{jk}(\mathbf{x}, \mathbf{y})\omega^{-k},$$

where the τ_j's are the traveltimes, each obeying (3.6) away from caustic curves. The amplitudes are singular at caustic curves in addition to the diagonal $\mathbf{x} = \mathbf{y}$, and contain the information of the Maslov indices. Note that traveltimes are symmetric: $\tau_j(\mathbf{x}, \mathbf{y}) = \tau_j(\mathbf{y}, \mathbf{x})$, and so is the kernel of D.

The singularity of the amplitude factor in (3.5), at $\mathbf{x} = \mathbf{y}$, is $O(\log|\mathbf{x} - \mathbf{y}|)$ in 2D and $O(|\mathbf{x} - \mathbf{y}|^{-1})$ in 3D. After differentiating twice to obtain H, the homogeneity on the diagonal becomes $O(|\mathbf{x} - \mathbf{y}|^{-2})$ in 2D and $O(|\mathbf{x} - \mathbf{y}|^{-3})$ in 3D. For the decay at infinity, the scalings are different and can be obtained from Fourier analysis of square root singularities; the kernel of H decays like $O\left(|\mathbf{x} - \mathbf{y}|^{-3/2}\right)$ in 2D, and $O\left(|\mathbf{x} - \mathbf{y}|^{-5/2}\right)$ in 3D. In between, the amplitude is smooth as long as the traveltime is single-valued.

As mentioned before, much more is known about DtN maps, such as boundedness and coercivity theorems. Again, we did not attempt to leverage these properties of D in the scheme presented here.

For all these reasons, we define the basis matrices B_j as follows. Assume τ is single-valued. In 1D, denote the tangential component of $\mathbf{x}$ by x, and similarly that of $\mathbf{y}$ by y, in coordinates local to each edge with $0 \leq x, y \leq 1$. Each block M of D relates to a couple of edges of the square domain. Let $j = (j_1, j_2)$ with j_1, j_2 nonnegative integers. The general forms that we consider are

$$\beta_j(x, y) = e^{i\omega\tau(x,y)}(h + |x - y|)^{-\frac{j_1}{\alpha}}(h + \theta(x, y))^{-\frac{j_2}{\alpha}}$$

and

$$\beta_j(x, y) = e^{i\omega\tau(x,y)}(h + |x - y|)^{-\frac{j_1}{\alpha}}(h + \theta(x, y))^{j_2},$$

where again h is the grid spacing of the FD scheme, and $\theta(x, y)$ is an adequate function of x and y that depends on the particular block of interest. The more favorable choices for θ are those that respect the singularities created at the vertices of the square; we

typically let $\theta(x, y) = \min(x + y, 2 - x - y)$. The parameter α can be taken to be equal to 2, a good choice in view of the numerics and in the light of the asymptotic behaviors on the diagonal and at infinity discussed earlier.

If several traveltimes are needed for geometrical reasons, then different sets of β_j are defined for each traveltime. (More about this in the next subsection.) The B_j are then obtained from the β_j by QR factorization within each block[2], where orthogonality is defined in the sense of the Frobenius inner product $\langle A, B\rangle = \mathrm{tr}(AB^T)$. This automatically sets the κ number of probing to 1.

In many of our test cases it appears that the "triangular" condition $j_1 + 2j_2 <$ *constant* works well. The number of couples (j_1, j_2) satisfying this relation will be p/T, where p is the number of basis matrices in the matrix probing algorithm and T is the number of distinct traveltimes. The eventual ordering of the basis matrices B_j respects the increase of $j_1 + 2j_2$.

3.2.2 More on traveltimes

Determining the traveltime(s) $\tau(\mathbf{x}, \mathbf{y})$ is the more "supervised" part of this method, but is needed to keep the number p of parameters small in the probing expansion. A few different scenarios can arise.

- In the case when $\nabla c(\mathbf{x})$ is perpendicular to a straight segment of the boundary, locally, then this segment is itself a ray and the waves can be labeled as interfacial, or "creeping". The direct traveltime between any two points $\mathbf{x}$ and $\mathbf{y}$ on this segment is then simply given by the line integral of $1/c(\mathbf{x})$. An infinite sequence of additional interfacial waves result from successive reflections at the endpoints of the segment, with traveltimes predicted as follows.

  We still consider the exterior problem for $[0, 1]^2$. We are interested in the traveltimes between points $\mathbf{x}, \mathbf{y}$ on the same side of $\partial\Omega$ – for illustration, let $\mathbf{x} = (x, 0)$ and $\mathbf{y} = (y, 0)$ on the bottom side of $\Omega = [0, 1]^2$, with $x \leq y$ (this is sufficient since traveltimes are symmetric). Assume that all the waves are interfacial. The first traveltime τ_1 corresponds to the direct path from $\mathbf{x}$ to $\mathbf{y}$. The second arrival time τ_2 will be the minimum traveltime corresponding to: either starting at $\mathbf{x}$, going left, reflecting off of the $(0, 0)$ corner, and coming back along the bottom side of $\partial\Omega$, past $\mathbf{x}$ to finally reach $\mathbf{y}$; or starting at $\mathbf{x}$, going past $\mathbf{y}$, reflecting off of the $(1, 0)$ and coming straight back to $\mathbf{y}$. The third arrival time τ_3 is the maximum of those two choices. The fourth arrival time then corresponds to starting at $\mathbf{x}$, going left, reflecting off of the $(0, 0)$ corner, travelling all the way to the $(1, 0)$ corner, and then back to $\mathbf{y}$. The fifth arrival time corresponds to leaving $\mathbf{x}$, going to the $(1, 0)$ corner this time, then back to the $(0, 0)$ corner,

[2]Whenever a block of D has symmetries, we enforce those in the QR factorization by using appropriate weights on a subset of the entries of that block. This also reduces the complexity of the QR factorization.

then on to $\mathbf{y}$. And so on. To recap, we have the following formulas:

$$\begin{aligned}
\tau_1(\mathbf{x},\mathbf{y}) &= \int_x^y \frac{1}{c(t,0)}\,dt,\\
\tau_2(\mathbf{x},\mathbf{y}) &= \tau_1(\mathbf{x},\mathbf{y}) + 2\min\left(\int_0^x \frac{1}{c(t,0)}\,dt, \int_y^1 \frac{1}{c(t,0)}\,dt\right),\\
\tau_3(\mathbf{x},\mathbf{y}) &= \tau_1(\mathbf{x},\mathbf{y}) + 2\max\left(\int_0^x \frac{1}{c(t,0)}\,dt, \int_y^1 \frac{1}{c(t,0)}\,dt\right) = 2\int_0^1 \frac{1}{c(t,0)}\,dt - \tau_2(\mathbf{x},\mathbf{y}),\\
\tau_4(\mathbf{x},\mathbf{y}) &= 2\int_0^1 \frac{1}{c(t,0)}\,dt - \tau_1(\mathbf{x},\mathbf{y}),\\
\tau_5(\mathbf{x},\mathbf{y}) &= 2\int_0^1 \frac{1}{c(t,0)}\,dt + \tau_1(\mathbf{x},\mathbf{y}), \qquad \text{etc.}
\end{aligned}$$

All first five traveltimes can be expressed as a sum of $\pm\tau_1$, $\pm\tau_2$ and the constant phase $2\int_0^1 \frac{1}{c(t,0)}\,dt$, which does not depend on $\mathbf{x}$ or $\mathbf{y}$. In fact, one can see that any subsequent traveltime corresponding to traveling solely along the bottom boundary of $\partial\Omega$ should be again a combination of those quantities. This means that if we use $\pm\tau_1$ and $\pm\tau_2$ in our basis matrices, we are capturing all the traveltimes relative to a single side, which helps to obtain higher accuracy for probing the diagonal blocks of D.

This simple analysis can be adapted to deal with creeping waves that start on one side of the square and terminate on another side, which is important for the nondiagonal blocks of D.

- In the case when $c(\mathbf{x})$ increases outward in a smooth fashion, we are also often in presence of body waves, going off into the exterior and coming back to $\partial\Omega$. The traveltime for these waves needs to be solved either by a Lagrangian method (solving the ODE for the rays), or by an Eulerian method (solving the Eikonal PDE shown earlier). In this paper we used the fast marching method of Sethian [46] to deal with these waves in the case that we label "slow disk" in the next section.

- In the case when $c(\mathbf{x})$ has singularities in the exterior domain, each additional reflection creates a traveltime that should (ideally) be predicted. Such is the case of the "diagonal fault" example introduced in the next section, where a straight jump discontinuity of $c(\mathbf{x})$ intersects $\partial\Omega$ at a non-normal angle: we can construct by hand the traveltime corresponding to a path leaving the boundary at $\mathbf{x}$, reflecting off of the discontinuity and coming back to the boundary at $\mathbf{y}$. More precisely, we consider again $\mathbf{x} = (x, 0)$, $\mathbf{y} = (y, 0)$ and $x \le y$, with x larger than or equal to the x coordinate of the point where the reflector intersects the bottom side of $\partial\Omega$. We then reflect the point $\mathbf{y}$ across the discontinuity into the new point $\mathbf{y}'$, and calculate the Euclidean distance between $\mathbf{x}$ and $\mathbf{y}'$. To obtain the traveltime, we then divide this distance by the value $c(\mathbf{x}) = c(\mathbf{y})$ of c on

the right side of the discontinuity, assuming that value is constant. This body traveltime is used in the case of the "diagonal fault", replacing the quantity τ_2 that was described above. This increased accuracy by an order of magnitude, as mentioned in the numerical results of the next section.

3.3 Numerical experiments

Our benchmark media $c(\mathbf{x})$ are as follows:

1. a uniform wave speed of 1, $c \equiv 1$ (Figure 3-1),
2. a "Gaussian waveguide" (Figure 3-2),
3. a "Gaussian slow disk" (Figure 3-3) large enough to encompass Ω – this will cause some waves going out of Ω to come back in,
4. a "vertical fault" (Figure 3-4),
5. a "diagonal fault" (Figure 3-5),
6. and a discontinuous periodic medium (Figure 3-6). The periodic medium consists of square holes of velocity 1 in a background of velocity $1/\sqrt{12}$.

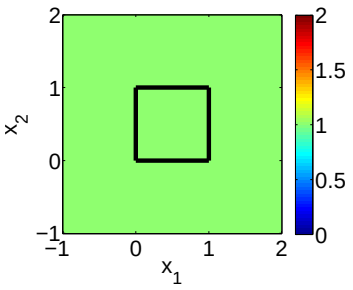

Figure 3-1: Color plot of $c(\mathbf{x})$ for the uniform medium.

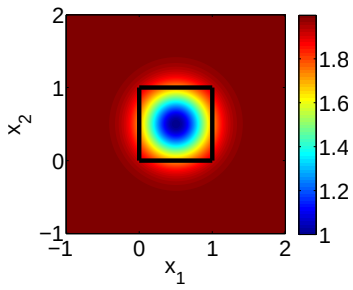

Figure 3-3: Color plot of $c(\mathbf{x})$ for the Gaussian slow disk.

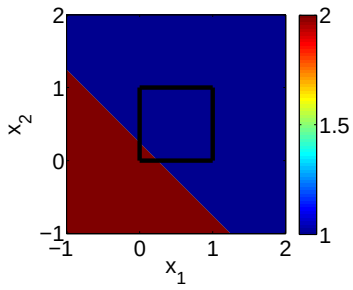

Figure 3-5: Color plot of $c(\mathbf{x})$ for the diagonal fault.

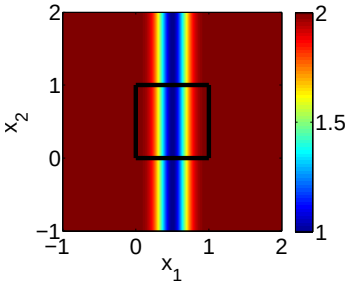

Figure 3-2: Color plot of $c(\mathbf{x})$ for the Gaussian waveguide.

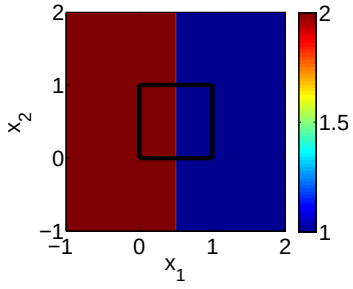

Figure 3-4: Color plot of $c(\mathbf{x})$ for the vertical fault.

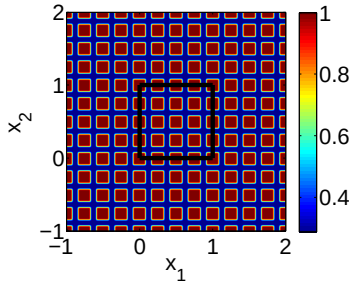

Figure 3-6: Color plot of $c(\mathbf{x})$ for the periodic medium.

Medium	N	$\omega/2\pi$	FD error	w	P	Source position
$c \equiv 1$	1023	51.2	$2.5e-01$	4	8	$(0.5, 0.25)$
waveguide	1023	51.2	$2.0e-01$	4	56	$(0.5, 0.5)$
slow disk	1023	51.2	$1.8e-01$	4	43	$(0.5, 0.25)$
fault, left source	1023	51.2	$1.1e-01$	4	48	$(0.25, 0.5)$
fault, right source	1023	51.2	$2.2e-01$	4	48	$(0.75, 0.5)$
diagonal fault	1023	51.2	$2.6e-01$	256	101	$(0.5, 0.5)$
periodic medium	319	6	$1.0e-01$	1280	792	$(0.5, 0.5)$

Table 3.1: For each medium considered, we show the parameters N and $\omega/2\pi$, along with the resulting discretization error caused by the Finite Difference (FD error) formulation. We also show the width w of the pPML needed, in number of points, to obtain an error caused by the pPML of less than $1e-1$. Furthermore, we show the total number P of basis matrices needed to probe the entire DtN map with an accuracy of about $1e-1$ as found in Section 3.3.1. Finally, we show the position of the point source used in calculating the solution u.

All media used are continued in the obvious way (i.e., they are *not* put to a uniform constant) outside of the domain in which they are shown in the figures if needed. The outline of the $[0,1]^2$ box is shown in black.

We can use a standard Helmholtz equation solver to estimate the relative error in the Helmholtz equation solution caused by the Finite Difference discretization (the *FD error*[3]), and also the error caused by using the specified pPML width[4]. Those errors are presented in Table 3.1, along with the main parameters used in the remaining of this section, including the position of the point source or right-hand side f. We note that, whenever possible, we try to use an AL with error smaller than the precision we seek with matrix probing, so with a width w greater than that showed in Table 3.1. This makes probing easier, i.e. p and q can be smaller.

Consider now a block M of D, corresponding to the restriction of D to two sides of $\partial\Omega$. We note that some blocks in D are the same up to transpositions or flips (inverting the order of columns or rows) if the medium c has symmetries.

Definition 6. Multiplicity of a block of D. *Let M be a block of D, corresponding to the restriction of D to two sides of $\partial\Omega$. The* multiplicity $m(M)$ *of M is the number of copies of M appearing in D, up to transpositions or flips.*

Only the distinct blocks of D need to be probed. Once we have chosen a block M, we may calculate the *true probing coefficients*.

Definition 7. True probing coefficients of block M. *Let M be a block of D, corresponding to the restriction of D to two sides of $\partial\Omega$. Assume orthonormal probing basis matrices $\{B_j\}$. The true coefficients c_j^t in the probing expansion of M are the inner products $c_j^t = \langle B_j, M \rangle$.*

[3]To find this FD error, we use a large pseudo-PML, and compare the solution u for different values of N. What we call the FD error is the relative ℓ_2 error in u inside Ω.

[4]To obtain the error caused by the absorbing layer, we fix N and compare the solution u for different layer widths w, and calculate the relative ℓ_2 error in u inside Ω.

We may now define the *p-term approximation error* for the block M.

Definition 8. The p-term approximation error of block M. *Let M be a block of D, corresponding to the restriction of D to two sides of $\partial\Omega$. For orthonormal probing basis matrices $\{B_j\}$, we have the true coefficients c_j^t in the probing expansion of M. Let $M_p = \sum_{j=1}^{p} c_j^t B_j$ be the probing p-term approximation to M. The p-term approximation error for M is*

$$\sqrt{m(M)}\frac{\|M - M_p\|_F}{\|D\|_F}, \tag{3.7}$$

using the matrix Frobenius norm.

Because the blocks on the diagonal of D have a singularity, their Frobenius norm can be a few orders of magnitude greater than that of other blocks, and so it is more important to approximate those well. This is why we consider the error relative to D, not to the block M, in the p-term approximation error. Also, we multiply by the square root of the multiplicity of the block to give us a better idea of how big the total error on D will be. For brevity, we shall refer to (3.7) simply as the approximation error when it is clear from the context what M, p $\{B_j\}$, D are.

Then, using matrix probing, we will recover a coefficient vector $\mathbf{c}$ close to $\mathbf{c}^t$, which gives an approximation $\widetilde{M} = \sum_{j=1}^{p} c_j B_j$ to M. We now define the *probing error* (which depends on q and the random vectors used), for the block M.

Definition 9. Probing error of block M. *Let $\mathbf{c}$ be the probing coefficients for M obtained with q random realizations $z^{(1)}$ through $z^{(q)}$. Let $\widetilde{M} = \sum_{j=1}^{p} c_j B_j$ be the probing approximation to M. The probing error of M is*

$$\sqrt{m(M)}\frac{\|M - \widetilde{M}\|_F}{\|D\|_F}. \tag{3.8}$$

Again, for brevity, we refer to (3.8) as the probing error when other parameters are clear from the context. Once all distinct blocks of D have been probed, we can consider the *total probing error*.

Definition 10. Total probing error. *The total probing error is defined as the total error made on D by concatenating all probed blocks $\widetilde{M}$ to produce an approximate $\widetilde{D}$, and is equal to*

$$\frac{\|D - \widetilde{D}\|_F}{\|D\|_F}. \tag{3.9}$$

In order to get a point of reference for the accuracy benchmarks, for small problems only, the actual matrix D is computed explicitly by solving the exterior problem $4N$ times using the standard basis as Dirichlet boundary conditions, and from this we can calculate (3.9) exactly. For larger problems, we only have access to a black-box that outputs the product of D with some input vector by solving the exterior problem. We can then estimate (3.9) by comparing the products of D and $\widetilde{D}$ with a few random vectors different from those used in matrix probing.

We shall present results on the approximation and probing errors for various media, along with related condition numbers, and then we shall verify that using an approximate $\widetilde{D}$ (constructed from approximate $\widetilde{M}$'s for each block M in D) does not affect the accuracy of the new solution to the Helmholtz equation, using the *solution error from probing*.

Definition 11. Solution error from probing. *Once we have obtained an approximation $\widetilde{D}$ to D from probing the distinct blocks of D, we may use this $\widetilde{D}$ in a Helmholtz solver to obtain an approximate solution $\widetilde{u}$, and compare that to the true solution u using D in the solver. The solution error from probing is the ℓ_2 error on u inside Ω:*

$$\frac{\|u - \widetilde{u}\|_2}{\|u\|_2} \text{ in } \Omega. \tag{3.10}$$

3.3.1 Probing tests

As we saw in Section 3.1.2, randomness plays a role in the value of cond($\mathbf{\Psi}$) and of the probing error. Hence, whenever we show plots for those quantities in this section, we have done 10 trials for each value of q used. The error bars show the minimum and maximum of the quantity over the 10 trials, and the line is plotted through the average value over the 10 trials. As expected, we will see in all experiments that increasing q gives a better conditioning, and consequently a better accuracy and smaller failure probability. The following probing results will then be used in Section 3.3.2 to solve the Helmholtz equation.

Uniform medium

For a uniform medium, $c \equiv 1$, we have three blocks with the following multiplicities: $m((1,1)) = 4$ (same edge), $m((2,1)) = 8$ (neighboring edges), and $m((3,1)) = 4$ (opposite edges). Note that we do not present results for the $(3,1)$ block: this block has negligible Frobenius norm[5] compared to D. First, let us look at the conditioning for blocks $(1,1)$ and $(2,1)$. Figures 3-7 and 3-8 show the three relevant conditioning quantities: κ, λ and cond($\mathbf{\Psi}$) for each block. As expected, $\kappa = 1$ because we orthogonalize the basis functions. Also, we see that λ does not grow very much as p increases, it remains on the order of 10. As for cond($\mathbf{\Psi}$), it increases as p increases for a fixed q and N, as expected. This will affect probing in terms of the failure probability (the odds that the matrix $\mathbf{\Psi}$ is far from the expected value) and accuracy (taking the pseudo-inverse will introduce larger errors in $\mathbf{c}$). We notice these two phenomena in Figure 3-9, where we show the approximation and probing errors in probing the $(1,1)$ block for various p, using different q and making 10 tests for each q value as explained previously. As expected, as p increases, the variations between trials get larger. Also, the probing error, always larger than the approximation error, becomes farther and farther away from the approximation error. Comparing Figure 3-9 with Table 3.2 of the next section, we see that in Table 3.2 we are able to achieve higher

[5] We can use probing with $q = 1$ and a single basis matrix (a constant multiplied by the correct oscillations) and have a probing error of less than 10^{-6} for that block.

accuracies. This is because we use the first two traveltimes (so four different types of oscillations, as explained in Section 3.2) to obtain those higher accuracies. But we do not use four types of oscillations for lower accuracies because this demands a larger number of basis matrices p and of solves q for the same error level.

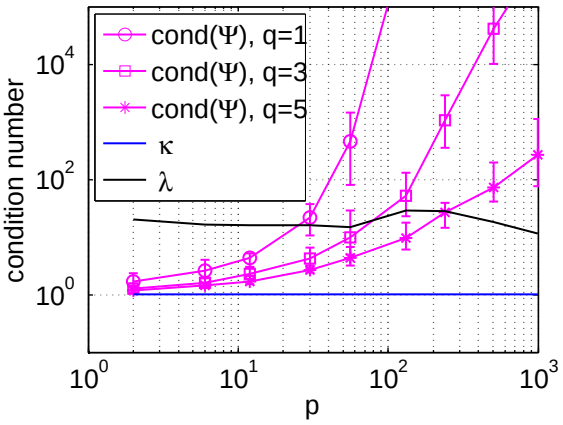

Figure 3-7: Condition numbers for the $(1,1)$ block, $c \equiv 1$.

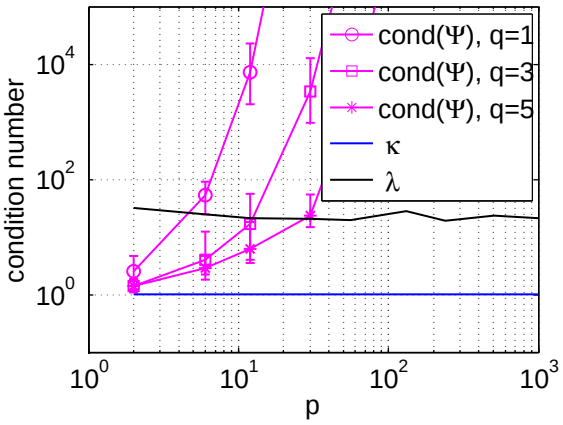

Figure 3-8: Condition numbers for the $(2,1)$ block, $c \equiv 1$.

The waveguide

For a waveguide as a velocity field, we have more blocks compared to the uniform medium case, with different multiplicities: $m((1,1)) = 2$, $m((2,2)) = 2$, $m((2,1)) = 8$, $m((3,1)) = 2$, $m((4,2)) = 2$. Note that block $(2,2)$ will be easier to probe than block $(1,1)$ since the medium is smoother on that interface. Also, we can probe blocks $(3,1)$ and $(4,2)$ with $q = 1$, $p = 2$ and have a probing error less than 10^{-7}. Hence we only show results for the probing and approximation errors of blocks $(1,1)$ and $(2,1)$, in Figure 3-10. Results for using probing in a solver can be found in Section 3.3.2.

The slow disk

Next, we consider the slow disk. Here, we have a choice to make for the traveltime upon which the oscillations depend. We may consider interfacial waves, traveling in straight line segments along $\partial\Omega$, with traveltime τ. There is also the first arrival time of body waves, τ_f, which for some points on $\partial\Omega$ involve taking a path that goes away from $\partial\Omega$, into the exterior where c is higher, and back towards $\partial\Omega$. We have approximated this τ_f using the fast marching method of Sethian [46]. For this example, it turns out that using either τ or τ_f to obtain oscillations in our basis matrices does not significantly alter the probing accuracy or conditioning, although it does seem that, for higher accuracies at least, the fast marching traveltime makes convergence slightly faster. Figures 3-11 and 3-12 demonstrate this for blocks $(1,1)$ and $(2,1)$ respectively. We omit plots of the probing and approximation errors, and refer the reader to Section 3.3.2 for final probing results and using those in a solver.

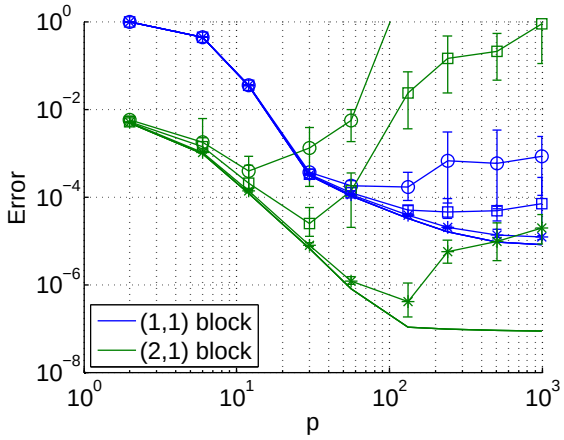

Figure 3-9: Approximation error (line) and probing error (with markers) for the blocks of D, $c \equiv 1$. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$.

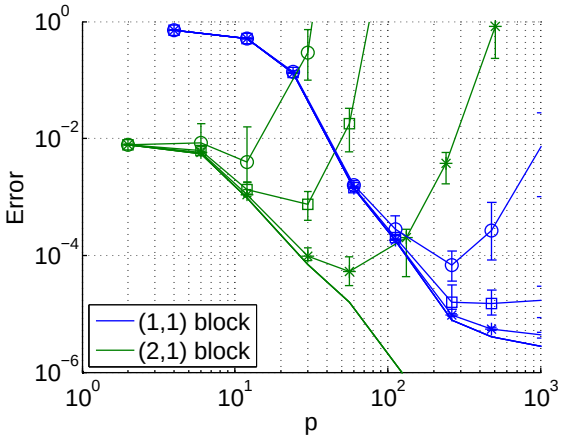

Figure 3-10: Approximation error (line) and probing error (with markers) for the blocks of D, c is the waveguide. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$.

The vertical fault

Next, we look at the case of the medium c which has a vertical fault. We note that this case is harder because some of the blocks will have themselves a 2 by 2 or 1 by 2 structure caused by the discontinuity in the medium. Ideally, as we shall see, each sub-block should be probed separately. There are 7 distinct blocks, with different multiplicities: $m((1,1)) = 2$, $m((2,2)) = 1$, $m((4,4)) = 1$, $m((2,1)) = 4$, $m((4,1)) = 4$, $m((3,1)) = 2$, $m((4,2)) = 2$. Blocks $(2,2)$ and $(4,4)$ are easier to probe than block $(1,1)$ because they do not exhibit a sub-structure. Also, since the velocity is smaller on the right side of the fault, the frequency there is higher, which means that blocks involving side 2 are slightly harder to probe than those involving side 4. Hence we first present results for the blocks $(1,1)$, $(2,2)$ and $(2,1)$ of D. In Figure 3-13 we see the approximation and probing errors for those blocks. Then, in Figure 3-14, we present results for the errors related to probing the 3 distinct sub-blocks of the $(1,1)$ block of D. We can see that probing the $(1,1)$ block by sub-blocks helps achieve greater accuracy. We could have split other blocks too to improve the accuracy of their probing (for example, block $(2,1)$ has a 1 by 2 structure because side 1 has a discontinuity in c) but the accuracy of the overall DtN map was still limited by the accuracy of probing the $(1,1)$ block, so we do not show results for other splittings.

The diagonal fault

Now, we look at the case of the medium c which has a diagonal fault. Again, some of the blocks will have themselves a 2 by 2 or 1 by 2 structure. There are 6 distinct blocks, with different multiplicities: $m((1,1)) = 2$, $m((2,2)) = 2$, $m((2,1)) = 4$, $m((4,1)) = 2$, $m((3,2)) = 2$, $m((3,1)) = 4$. Again, we split up block $(1,1)$ in 4 sub-blocks and probe each of those sub-blocks separately for greater accuracy, but

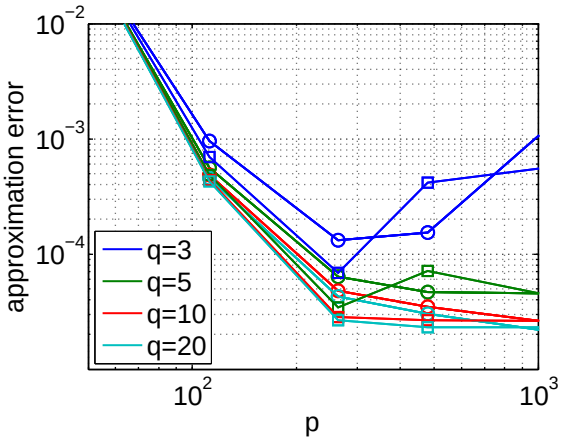

Figure 3-11: Approximation error for the $(1,1)$ blocks of D, c is the slowness disk, comparing the use of the normal traveltime (circles) to the fast marching traveltime (squares).

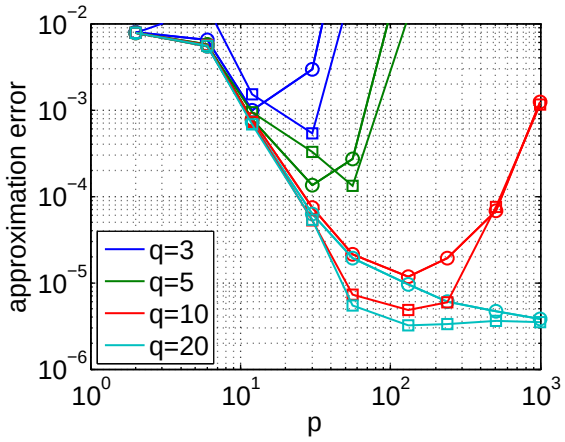

Figure 3-12: Approximation error for the $(2,1)$ blocks of D, c is the slowness disk, comparing the use of the normal traveltime (circles) to the fast marching traveltime (squares).

do not split other blocks. We then use two traveltimes for the $(2,2)$ sub-block of block $(1,1)$. Using as the second arrival time the geometrical traveltime consisting of leaving the boundary and bouncing off the fault, as mentioned in Section 3.2.2, allowed us to increase accuracy by an order of magnitude compared to using only the first arrival traveltime, or compared to using as a second arrival time the usual bounce off the corner (or here, bounce off the fault where it meets $\delta\Omega$). We omit plots of the probing and approximation errors, and refer the reader to Section 3.3.2 for final probing results and using those in a solver.

The periodic medium

Finally, we look at the case of the periodic medium presented earlier. There are 3 distinct blocks, with different multiplicities: $m((1,1)) = 4$, $m((2,1)) = 8$, $m((3,1)) = 4$. We expect the corresponding DtN map to be harder to probe because its structure will reflect that of the medium, i.e. it will exhibit sharp transitions at points corresponding to sharp transitions in c (similarly as with the faults). First, we notice that, in all the previous mediums we tried, plotting the norm of the anti-diagonal entries of diagonal blocks (or sub-blocks for the faults) shows a rather smooth decay away from the diagonal. However, that is not the case for the periodic medium: it looks like there is decay away from the diagonal, but variations from that decay can be of relative order 1. This prevents our usual strategy, using basis matrices containing terms that decay away from the diagonal such as $(h + |x - y|)^{-j_1/\alpha}$, from working adequately. Instead, we use polynomials along anti-diagonals, as well as polynomials along diagonals as we previously did.

It is known that solutions to the Helmholtz equation in a periodic medium are Bloch waves with a particular structure [33]. However, using that structure in the

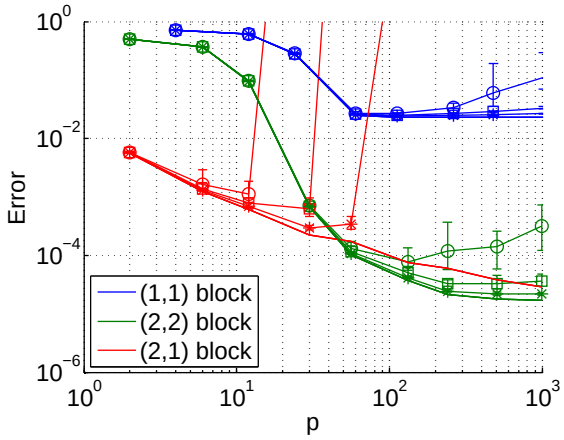

Figure 3-13: Approximation error (line) and probing error (with markers) for the blocks of D, c is the fault. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$.

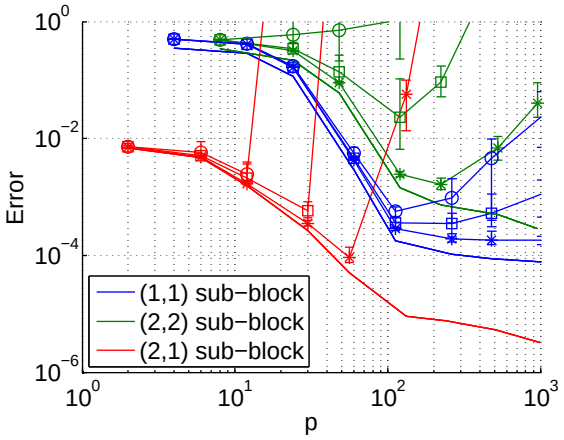

Figure 3-14: Approximation error (line) and probing error (with markers) for the sub-blocks of the $(1, 1)$ block of D, c is the fault. Circles are for $q = 3$, squares for $q = 5$, stars for $q = 10$.

basis matrices is not robust. Indeed, using a Bloch wave structure did not succeed very well, probably because our discretization was not accurate enough and so D exhibited that structure only to a very rough degree. Hence we did not use Bloch waves for probing the periodic medium. Others have successfully used the known structure of the solution in this setting to approximate the DtN map. In [23], the authors solve local cell problems and Ricatti equations to obtain discrete DtN operators for media which are a perturbation of a periodic structure. In [55], the authors develop a DtN map eigenvalue formulation for wave propagation in periodic media. We did not attempt to use those formulations here.

For this reason, we tried basis matrices with no oscillations, but with polynomials in both directions as explained previously, and obtained the results of Section 3.3.2.

Now that we have probed the DtN map and obtained compressed blocks to form an approximation $\widetilde{D}$ of D, we may use this $\widetilde{D}$ in a Helmholtz solver as an absorbing boundary condition.

3.3.2 Using the probed DtN map in a Helmholtz solver

In Figures 3-15, 3-16, 3-17, 3-18, 3-19 and 3-20 we can see the standard solutions to the Helmholtz equation on $[0, 1]^2$ using a large PML or pPML for the various media we consider, except for the uniform medium, where the solution is well-known. We use those as our reference solutions.

We have tested the solver with the probed $\widetilde{D}$ as an absorbing boundary condition with success. See Tables 3.2, 3.4, 3.3 and 3.5 for results corresponding to each medium. We show the number p of basis matrices required for some blocks for that tolerance, the number of solves q of the exterior problem for those blocks, the total number of solves Q, the total probing error (3.9) in D and the solution error from probing

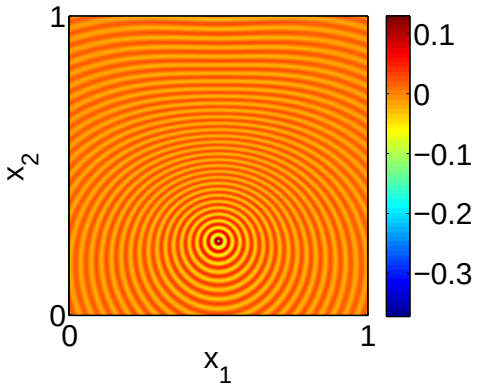

Figure 3-15: Real part of the solution, c is the slow disk.

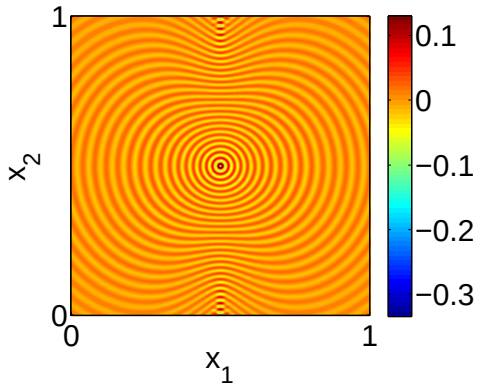

Figure 3-16: Real part of the solution, c is the waveguide.

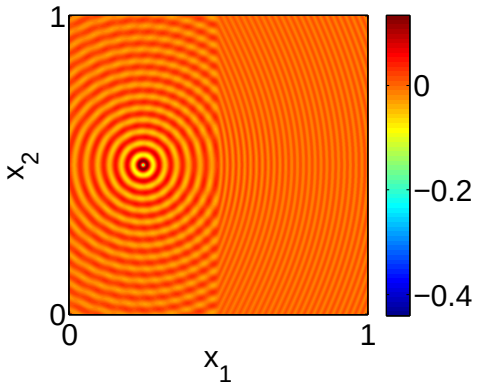

Figure 3-17: Real part of the solution, c is the vertical fault with source on the left.

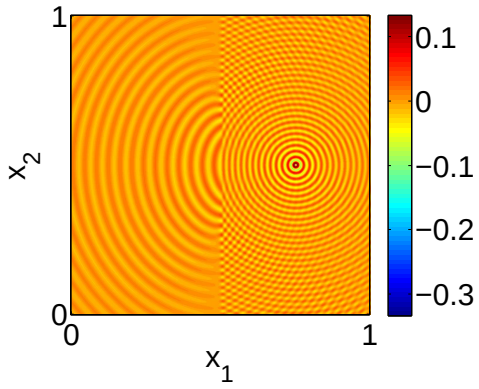

Figure 3-18: Real part of the solution, c is the vertical fault with source on the right.

(3.10). As we can see from the tables, the solution error from probing (3.10) in the solution u is no more than an order of magnitude greater than the total probing error (3.9) in the DtN map D, for a source position as described in Table 3.1. Grazing waves, which can arise when the source is close to the boundary of the computational domain, will be discussed in the next subsection, 3.3.3. We note again that, for the uniform medium, using the second arrival traveltime as well as the first for the $(1,1)$ block allowed us to achieve accuracies of 5 and 6 digits in the DtN map, which was not possible otherwise. Using a second arrival time for the cases of the faults was also useful. Those results show that probing works best when the medium c is rather smooth. For non-smooth media such as a fault, it becomes harder to probe the DtN map to a good accuracy, so that the solution to the Helmholtz equation also contains more error.

Table 3.2: $c \equiv 1$

p for $(1,1)$	p for $(2,1)$	$q = Q$	$\frac{\|D-\widetilde{D}\|_F}{\|D\|_F}$	$\frac{\|u-\widetilde{u}\|_2}{\|u\|_2}$
6	1	1	$2.0130e-01$	$3.3191e-01$
12	2	1	$9.9407e-03$	$1.9767e-02$
20	12	3	$6.6869e-04$	$1.5236e-03$
72	20	5	$1.0460e-04$	$5.3040e-04$
224	30	10	$8.2892e-06$	$9.6205e-06$
360	90	10	$7.1586e-07$	$1.3044e-06$

Table 3.3: c is the waveguide

p for $(1,1)$	p for $(2,1)$	q	p for $(2,2)$	q	Q	$\frac{\|D-\widetilde{D}\|_F}{\|D\|_F}$	$\frac{\|u-\widetilde{u}\|_2}{\|u\|_2}$
40	2	1	12	1	2	$9.1087e-02$	$1.2215e-01$
40	2	3	20	1	4	$1.8685e-02$	$7.6840e-02$
60	20	5	20	3	8	$2.0404e-03$	$1.3322e-02$
112	30	10	30	3	13	$2.3622e-04$	$1.3980e-03$
264	72	20	168	10	30	$1.6156e-05$	$8.9911e-05$
1012	240	20	360	10	30	$3.3473e-06$	$1.7897e-05$

Table 3.4: c is the slow disk

p for $(1,1)$	p for $(2,1)$	$q = Q$	$\frac{\|D-\widetilde{D}\|_F}{\|D\|_F}$	$\frac{\|u-\widetilde{u}\|_2}{\|u\|_2}$
40	2	3	$1.0730e-01$	$5.9283e-01$
84	2	3	$8.0607e-03$	$4.5735e-02$
180	12	3	$1.2215e-03$	$1.3204e-02$
264	30	5	$1.5073e-04$	$7.5582e-04$
1012	132	20	$2.3635e-05$	$1.5490e-04$

Table 3.5: c is the fault

Q	$\frac{\|D-\widetilde{D}\|_F}{\|D\|_F}$	$\frac{\|u-\widetilde{u}\|_2}{\|u\|_2}$, left source	$\frac{\|u-\widetilde{u}\|_2}{\|u\|_2}$, right source
5	$2.8376e-01$	$6.6053e-01$	$5.5522e-01$
5	$8.2377e-03$	$3.8294e-02$	$2.4558e-02$
30	$1.1793e-03$	$4.0372e-03$	$2.9632e-03$

Table 3.6: c is the diagonal fault

Q	$\frac{\|D-\widetilde{D}\|_F}{\|D\|_F}$	$\frac{\|u-\widetilde{u}\|_2}{\|u\|_2}$
4	$1.6030e-01$	$4.3117e-01$
6	$1.7845e-02$	$7.1500e-02$
23	$4.2766e-03$	$1.2429e-02$

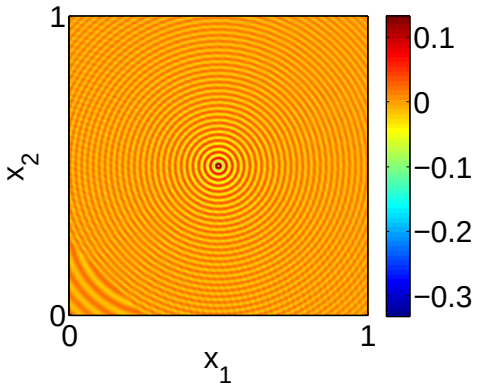

Figure 3-19: Real part of the solution, c is the diagonal fault.

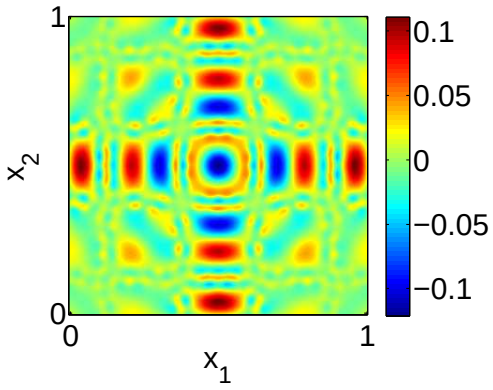

Figure 3-20: Imaginary part of the solution, c is the periodic medium.

Table 3.7: c is the periodic medium

Q	$\frac{\|D-\widetilde{D}\|_F}{\|D\|_F}$	$\frac{\|u-\widetilde{u}\|_2}{\|u\|_2}$
50	$1.8087e-01$	$1.7337e-01$
50	$3.5714e-02$	$7.1720e-02$
50	$9.0505e-03$	$2.0105e-02$

3.3.3 Grazing waves

It is well-known that ABCs often have difficulties when a source is close to a boundary of the domain, or in general when waves incident to the boundary are almost parallel to it. We wish to verify that the solution $\widetilde{u}$ using the result $\widetilde{D}$ of probing D does not degrade as the source becomes closer and closer to some side of $\partial\Omega$. For this, we use a right-hand side f to the Helmholtz equation which is a point source, located at the point (x_0, y_0), where $x_0 = 0.5$ is fixed and $y_0 > 0$ becomes smaller and smaller, until it is a distance $2h$ away from the boundary (the point source's stencil has width h, so a source at a distance h from the boundary does not make sense). We see in Figure 3-21 that, for $c \equiv 1$, the solution remains quite good until the source is a distance $2h$ away from the boundary. In this figure, we have used the probed maps we obtained in each row of Table 3.2. We obtain very similar results for the waveguide, slow disk and faults (for the vertical fault we locate the source at (x_0, y_0), where $y_0 = 0.5$ is fixed and x_0 goes to 0 or 1). This shows that the probing process itself does not significantly affect how well grazing waves are absorbed.

3.3.4 Variations of p with N

We now discuss how the number of basis matrices p needed to achieve a desired accuracy depends on N or ω. To do this, we pick 4 consecutive powers of 2 as values for N, and find the appropriate ω such that the finite discretization error remains constant at 10^{-1}, so that in fact $N \sim \omega^{1.5}$ as we have previously mentioned. We then

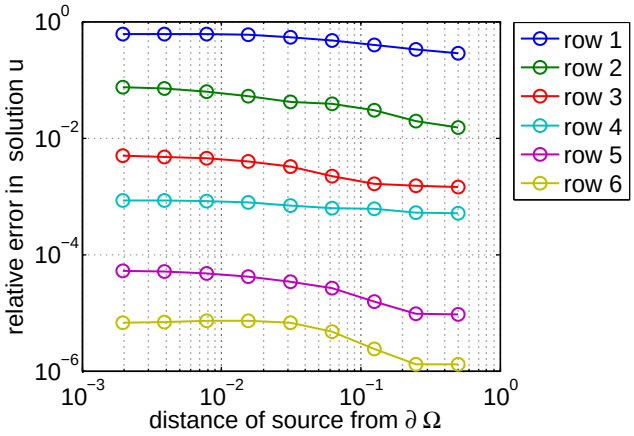

Figure 3-21: Error in solution u, $c \equiv 1$, moving point source. Each line in the plot corresponds to using $\widetilde{D}$ from a different row of Table 3.2.

probe the $(1,1)$ block of the corresponding DtN map, using the same parameters for all N, and observe the required p to obtain a fixed probing error. The worst case we have seen in our experiments came from the slow disk. As we can see in Figure 3-22, p seems to follow a very weak power law with N, close to $p \sim 15N^{.12}$ for a probing error of 10^{-1} or $p \sim 15N^{.2}$ for an probing error of 10^{-2}. In all other cases, p is approximately constant with increasing N, or seems to follow a logarithmic law with N as for the waveguide (see Figure 3-23).

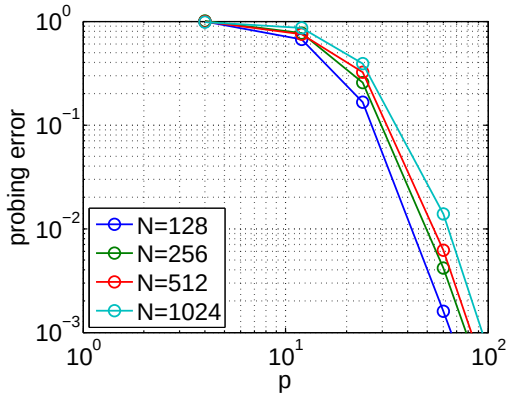

Figure 3-22: Probing error of the $(1,1)$ block of the DtN map for the slow disk, fixed FD error level of 10^{-1}, increasing N. This is the worst case, where p follows a weak power law with N.

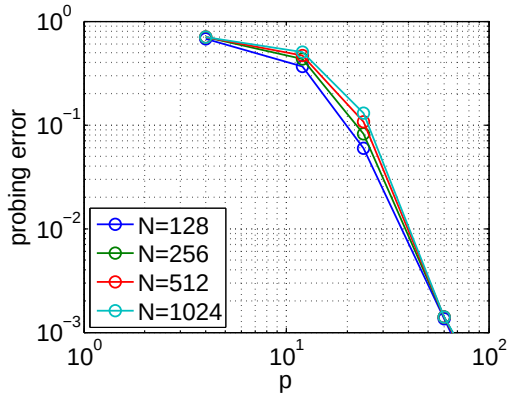

Figure 3-23: Probing error of the $(1,1)$ block of the DtN map for the waveguide, fixed FD error level of 10^{-1}, increasing N. Here p follows a logarithmic law.

3.4 Convergence of probing for the half-space DtN map: theorem

In this section, we consider the half-space DtN map kernel in uniform medium $K(r) = \frac{ik}{2r}H_1^{(1)}(kr)$ that we found in section 2.3.3. We wish to approximate this kernel for values of r that are relevant to our numerical scheme. Because we take Ω in our numerical experiments to be the $[0,1]^2$ box, $r = |x-y|$ will be between 0 and 1, in increments of h, as coordinates x and y along edges of $\partial\Omega$ vary between 0 and 1 in increments of h. However, as we know, $K(r)$ is singular at $r = 0$, and since discretization effects dominate near the diagonal in the matrix representation of the DtN map, we shall consider only values of r in the range $r_0 \leq r \leq 1$, with $0 < r_0 \leq 1/k$ (hence r_0 can be on the order of h). Since we know to expect oscillations e^{ikr} in this kernel, we can remove those from K to obtain

$$H(r) = \frac{ik}{2r}H_1^{(1)}(kr)e^{-ikr} \tag{3.11}$$

(not to be confused with the hypersingular kernel H of section 3.2), a smoother function which will be easier to approximate. Equivalently, we can add those oscillations to the terms in an approximation of H, to obtain an approximation of K.

For this section only, we denote by $\widetilde{D}$ the corresponding operator with integral kernel H, while we use D for the half-space Dirichlet-to-Neumann map, that is, the operator with kernel K.

Theorem 2. *Let $\alpha > \frac{2}{3}$, $0 < r_0 < 1/k$, and let $K_p(r)$ be the best uniform approximation of $K(r)$ in*

$$span\{\frac{e^{ikr}}{r^{j/\alpha}} : j = 1, \ldots, p, \text{ and } r_0 \leq r \leq 1\}.$$

Denote by D_p the operator defined with K_p in place of K. Then, in the operator norm,

$$\|D - D_p\| \leq C_\alpha\, p^{1-\lfloor 3\alpha/2 \rfloor}\, \|\widetilde{K}\|_\infty,$$

for some $C_\alpha > 0$ depending on α and $r_0 \leq r \leq 1$.

The important point of the theorem is that the quality of approximation is otherwise independent of k, i.e., the number p of basis functions does not need to grow like k for the error to be small. In other words, it is unnecessary to "mesh at the wavelength level" to spell out the degrees of freedom that go in the representation of the DtN map's kernel.

Remark 3.4.1. *Growing α does not automatically result in a better approximation error, because a careful analysis of the proof shows that C_α grows factorially with α. This behavior translates into a slower onset of convergence in p when α is taken large, as the numerics show in the next section. This can in turn be interpreted as the result of "overcrowding" of the basis by very look-alike functions.*

Remark 3.4.2. *It is easy to see that the operator norm of D grows like k, for instance by applying D to the function $e^{-ik\mathbf{x}}$. The uniform norms of K and H once we cut*

out the diagonal, however, grow like $k^{1/2}/r_0^{3/2}$, so the result above shows that we incur an additional factor $k^{-1/2}r_0^{-3/2}$ in the error (somewhat akin to numerical pollution) in addition to the factor k that we would have gotten from $\|D\|$.

The result in Theorem 2 points the way for the design of basis matrices to be used in matrix probing, for the more general case of the exterior DtN map in heterogeneous media. We prove 2 in the next subsections, and present a numerical verification in the next section.

3.4.3 Chebyshev expansion

We mentioned the domain of interest for the r variable is $[r_0, 1]$. Again, expanding $K(r)$ in the system of Theorem 2 is equivalent to expanding $H(r)$ in polynomials of $r^{-1/\alpha}$ over $[r_0, 1]$. It will be useful to perform the affine rescaling

$$\xi(r) = \frac{2}{r_0^{-1/\alpha} - 1}(r^{-1/\alpha} - 1) - 1 \qquad \Leftrightarrow \qquad r(\xi) = \left(\frac{\xi + 1}{2}(r_0^{-1/\alpha} - 1) + 1\right)^{-\alpha}$$

so that the bounds $r \in [r_0, 1]$ turn into $\xi \in [-1, 1]$. We further write $\xi = \cos\theta$ with $\theta \in [0, \pi]$. Our strategy is to expand H in Chebyshev polynomials $T_n(\xi)$. By definition, the best p-term approximation of $H(r)$ in polynomials of $r^{-1/\alpha}$ (best in a uniform sense over $[r_0, 1]$) will result in a lower uniform approximation error than that associated with the p-term approximation of $H(r(\xi))$ in the $T_n(\xi)$ system. Hence in the sequel we overload notations and let H_p for the p-term approximant of H in our Chebyshev system.

We write out the Chebyshev series for $H(r(\xi))$ as

$$H(r(\xi)) = \sum_{j=0}^{\infty} c_j T_j(\xi), \qquad c_j = \frac{2}{\pi}\int_{-1}^{1} \frac{H(r(\xi))T_j(\xi)}{(1-\xi^2)^{1/2}}\, d\xi,$$

with $T_j(\xi) = \cos\left(j(\cos^{-1}\xi)\right)$, and c_j alternatively written as

$$c_j = \frac{2}{\pi}\int_0^{\pi} H(r(\cos\theta))\cos j\theta\, d\theta = \frac{1}{\pi}\int_0^{2\pi} H(r(\cos\theta))\cos j\theta\, d\theta.$$

The expansion will converge fast because we can integrate by parts in θ and afford to take a few derivatives of H, say M of them, as done in [50]. After noting that the boundary terms cancel out because of periodicity in θ, we express the coefficients c_j for $j > 0$, up to a sign, as

$$c_j = \pm\frac{1}{\pi j^M}\int_0^{2\pi} \sin j\theta \frac{d^M}{d\theta^M} H(r(\cos\theta))\, d\theta, \qquad M \text{ odd},$$

$$c_j = \pm\frac{1}{\pi j^M}\int_0^{2\pi} \cos j\theta \frac{d^M}{d\theta^M} H(r(\cos\theta))\, d\theta, \qquad M \text{ even}.$$

It follows that, for $j > 0$, and for all $M > 0$,

$$|c_j| \leq \frac{2}{j^M} \max_\theta \left| \frac{d^M}{d\theta^M} H(r(\cos\theta)) \right|.$$

Let B_M be a bound on this M-th order derivative. The uniform error we make by truncating the Chebyshev series to $H_p = \sum_{j=0}^p c_j T_j$ is then bounded by

$$\|H - H_p\|_{L^\infty[r_0,1]} \leq \sum_{j=p+1}^\infty |c_j| \leq 2B_M \sum_{j=p+1}^\infty \frac{1}{j^M} \leq \frac{2B_M}{(M-1)p^{M-1}}, \qquad p > 1. \quad (3.12)$$

The final step is a simple integral comparison test.

3.4.4 Bound on the derivatives of the DtN map kernel with oscillations removed

The question is now to find a favorable estimate for B_M, from studying successive θ derivatives of $H(r)$ in (3.11). From the bound for the derivatives of Hankel functions in Lemma 1 of [16]: given any $C > 0$, we have

$$\left| \frac{d^m}{dr^m} \left(H_1^{(1)}(kr) e^{-ikr} \right) \right| \leq C_m (kr)^{-1/2} r^{-m} \qquad \text{for } kr \geq C. \quad (3.13)$$

The change of variables from r to θ results in

$$\begin{aligned} \frac{dr}{d\theta} &= \frac{d\xi}{d\theta}\frac{dr}{d\xi} = (-\sin\theta) \left(-\alpha \left(\frac{\xi+1}{2}(r_0^{-1/\alpha} - 1) + 1 \right)^{-\alpha-1} \frac{\left(r_0^{-1/\alpha} - 1 \right)}{2} \right) \\ &= (-\sin\theta) \left(-\alpha\ r^{1+1/\alpha}\ \frac{r_0^{-1/\alpha}(1 - r_0^{1/\alpha})}{2} \right). \end{aligned}$$

Hence

$$\frac{dr}{d\theta} = r(r/r_0)^{1/\alpha}\ \frac{\alpha \sin\theta (1 - r_0^{1/\alpha})}{2}. \quad (3.14)$$

Derivatives of higher powers of r are handled by the chain rule, resulting in

$$\frac{d}{d\theta}(r^p) = p r^p (r/r_0)^{1/\alpha}\ \frac{\alpha \sin\theta (1 - r_0^{1/\alpha})}{2}. \quad (3.15)$$

We see that the action of a θ derivative is essentially equivalent to multiplication by $(r/r_0)^{1/\alpha}$. As for higher derivatives of powers of r, it is easy to see by induction that the product rule has them either hit a power of r, or a trigonometric polynomial of θ, resulting in a growth of at most $(r/r_0)^{1/\alpha}$ for each derivative:

$$\left| \frac{d^m}{d\theta^m} r^p \right| \leq C_{m,p,\alpha}\, r^p (r/r_0)^{m/\alpha}.$$

These estimates can now be combined to bound $\frac{d^m}{d\theta^m}\left(H_1^{(1)}(kr)e^{-ikr}\right)$. One of two scenarios occur when applying the product rule:

- either $\frac{d}{d\theta}$ hits $\frac{d^{m_2}}{d\theta^{m_2}}\left(H_1^{(1)}(kr)e^{-ikr}\right)$ for some $m_2 < m$. In this case, one negative power of r results from $\frac{d}{dr}$ as we saw in (3.13), and a factor $r(r/r_0)^{1/\alpha}$ results from $\frac{dr}{d\theta}$ as we saw in (3.14);
- or $\frac{d}{d\theta}$ hits some power of r, possibly multiplied by some trigonometric polynomial in θ, resulting in a growth of an additional factor $(r/r_0)^{1/\alpha}$ as we saw in (3.15).

Thus, we get at most a $(r/r_0)^{1/\alpha}$ growth factor per derivative in every case. The situation is completely analogous when dealing with the slightly more complex expression $\frac{d^m}{d\theta^m}\left(\frac{1}{r}H_1^{(1)}(kr)e^{-ikr}\right)$. The number of terms is itself at most factorial in m, hence we get

$$|\frac{d^m}{d\theta^m}\frac{k}{r}\left(H_1^{(1)}(kr)e^{-ikr}\right)| \leq C_{m,\alpha}\ \frac{k}{r}\left(\frac{r}{r_0}\right)^{\frac{m}{\alpha}-\frac{1}{2}} \leq C_{m,\alpha}\ \frac{k}{r_0}\left(\frac{r}{r_0}\right)^{\frac{m}{\alpha}-\frac{3}{2}}. \qquad (3.16)$$

We now pick $m \leq M = \lfloor 3\alpha/2 \rfloor$, so that the max over θ is realized when $r = r_0$, and B_M is on the order of k/r_0. It follows from (3.12) and (3.16) that

$$\|H - H_p\|_{L^\infty[r_0,1]} \leq C_\alpha\ \frac{k}{r_0}\ \frac{1}{p^{\lfloor 3\alpha/2 \rfloor - 1}}, \qquad p > 1,\ \alpha > 2/3.$$

The kernel of interest, $K(r) = H(r)e^{ikr}$ obeys the same estimate if we let K_p be the p-term approximation of K in the Chebyshev system modulated by e^{ikr}.

3.4.5 Bound on the error of approximation

For ease of writing, we now let D^0, D_p^0 be the operators with respective kernels $K^0(r) = K(r)\chi_{[r_0,1]}(r)$ and $K_p^0(r) = K_p(r)\chi_{[r_0,1]}(r)$. We now turn to the operator norm of $D^0 - D_p^0$ with kernel $K^0 - K_p^0$:

$$(D^0 - D_p^0)g(x) = \int_0^1 (K^0 - K_p^0)(|x-y|)g(y)\ dy, \qquad x \in [0,1].$$

We use the Cauchy-Schwarz inequality to bound

$$\begin{aligned}
\|(D^0 - D_p^0)g\|_2 &= \left(\int_{0\leq x\leq 1}\left|\int_{0\leq y\leq 1,\ |x-y|\geq r_0}(K^0 - K_p^0)(|x-y|)g(y)\ dy\right|^2 dx\right)^{1/2} \\
&\leq \left(\int_{0\leq x\leq 1}\int_{0\leq y\leq 1,\ |x-y|\geq r_0}\left|(K^0 - K_p^0)(|x-y|)\right|^2\ dydx\right)^{1/2}\|g\|_2 \\
&\leq \left(\int_{0\leq x\leq 1}\int_{0\leq y\leq 1,\ |x-y|\geq r_0} 1\ dy\ dx\right)^{1/2}\|g\|_2 \max_{0\leq x,y\leq 1,\ |x-y|\geq r_0}|(K^0 - K_p^0)(|x-y|)| \\
&\leq \|g\|_2\ \|K^0 - K_p^0\|_{L^\infty[r_0,1]}.
\end{aligned}$$

Assembling the bounds, we have

$$\|D^0 - D_p^0\|_2 \leq \|K^0 - K_p^0\|_{L^\infty[r_0,1]} \leq C_\alpha \, p^{1-\lfloor 3\alpha/2 \rfloor} \, \frac{k}{r_0}.$$

It suffices therefore to show that $\|K^0\|_\infty = \|K\|_{L^\infty[r_0,1]}$ is bigger than k/r_0 to complete the proof. Letting $z = kr$, we see that

$$\max_{r_0 \leq r \leq 1} |K(r)| = \frac{k}{2r_0} \max_{kr_0 \leq z \leq k} \left| H_1^{(1)}(z) \right| \geq C \frac{k^{1/2}}{r_0^{3/2}}.$$

The last inequality follows from the fact that there exist a positive constant c_1 such that $c_1 z^{-1/2} \leq \left| H_1^{(1)}(z) \right|$, from Lemma 3 of [16]. But $k^{1/2}/r_0^{3/2} \geq k/r_0$ precisely when $r_0 \leq 1/k$. Hence we have proved the statement of Theorem 2.

In the next section, we proceed to a numerical confirmation of Theorem 2.

3.5 Convergence of probing for the half-space DtN map: numerical confirmation

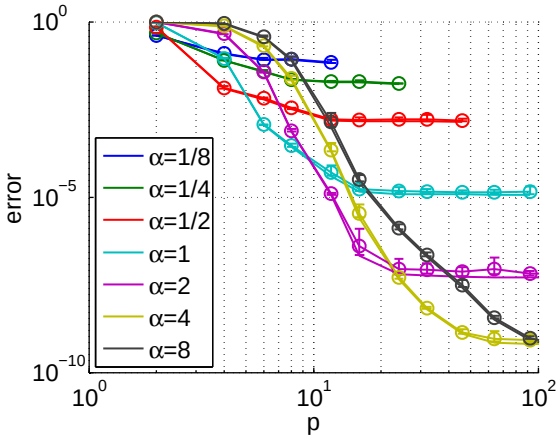

Figure 3-24: Probing error of the half-space DtN map($q = 1$, 10 trials, circle markers and error bars) compared to the approximation error (line), $c \equiv 1$, $L = 1/4$, $\alpha = 2$, $n = 1024$, $\omega = 51.2$.

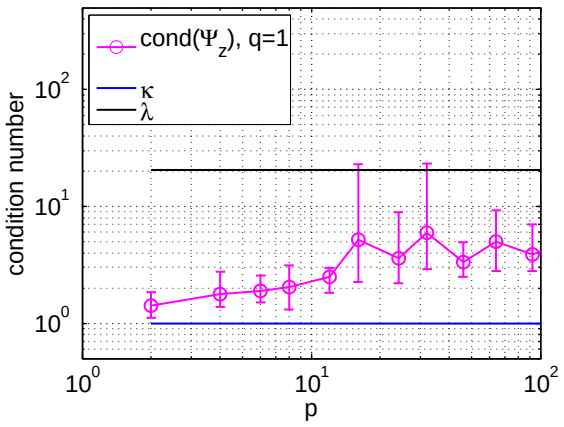

Figure 3-25: Condition numbers for probing the half-space DtN map, $c \equiv 1$, $L = 1/4$, $\alpha = 2$, $n = 1024$, $\omega = 51.2$, $q = 1$, 10 trials.

In order to use Theorem 2 to obtain convergent basis matrices, we start with the set $\left\{(r)^{-j/\alpha}\right\}_{j=0}^{p-1}$. We have proved the theorem for the interval $r_0 \leq r \leq 1$, but here we consider $h \leq r \leq 1$, which is a larger interval than in the theorem. We then put in oscillations, orthonormalize, and use this new set as a basis for probing the DtN

map. Thus we have pre-basis matrices ($0 \leq j \leq p$)

$$(\beta_j)_{\ell m} = \frac{e^{ikh|\ell - m|}}{|\ell - m|^{j/\alpha}} \text{ for } \ell \neq m,$$

with $(\beta_j)_{\ell\ell} = 0$. We add to this set the identity matrix in order to capture the diagonal of D, and orthonormalize the resulting collection to get the B_j. Alternatively, we have noticed that orthonormalizing the set of β_j's with

$$(\beta_j)_{\ell m} = \frac{e^{ikh|\ell - m|}}{(h + h|\ell - m|)^{j/\alpha}} \tag{3.17}$$

works just as well, and is simpler because there is no need to treat the diagonal separately. We use this same technique for the probing basis matrices of the exterior problem.

The convergent basis matrices in (3.17) have been used to obtain a numerical confirmation of Theorem 2, again for the half-space DtN map. To obtain the DtN map in this setup, instead of solving the exterior problem with a PML or pPML on all sides, we solve a problem on a thin strip, with a random Dirichlet boundary condition (for probing) one of the long edges, and a PML or pPML on the other three sides. This method for a numerical approximation of the solution to the half-space problem was introduced in section 2.3.2 and Figure 2-2.

3.5.1 Uniform medium

In Figure 3-24, we show the approximation error, which we expect will behave as in Theorem 2. We also plot error bars for the probing error, corresponding to ten trials of probing, with $q = 1$. The probing results are about as good as the approximation error, because the relevant condition numbers are all well-behaved as we see in Figure 3-25 for the value of choice of $\alpha = 2$. Back to the approximation error, we notice in Figure 3-24 that increasing α delays the onset of convergence as expected, because of the factor which is factorial in α in the error of Theorem 2. And we can see that, for small α, we are taking very high inverse powers of r, an ill-conditioned operation. Hence the appearance of a convergence plateau for smaller α is explained by ill-conditioning of the basis matrices, and the absence of data points is because of computational overflow.

Finally, increasing α from 1/8 to 2 gives a higher rate of convergence, as it should because in the error we have the factor $p^{-3\alpha/2}$, which gives a rate of convergence of $3\alpha/2$. This is roughly what we obtain numerically. As discussed, further increasing α is not necessarily advantageous since the constant C_α in 2 grows fast in α.

Chapter 4

Partitioned low-rank matrices for compressing the Dirichlet-to-Neumann map

In the previous chapter, we explained in detail the first step of two in our numerical scheme for compressing ABC's. This first step consisted in approximating the DtN map D by $\widetilde{D}$ using matrix probing. To do this, we considered each block M of D separately, corresponding to an edge-to-edge restriction of D. We approximated each M by a matrix probing expansion. We saw how we could obtain an accurate approximation of M using appropriate basis matrices B_j. We were left with the task of producing a fast algorithm for applying the resulting $\widetilde{M}$ to vectors, since this is an operation that a Helmholtz solver needs. This is what we do in the current chapter, by compressing $\widetilde{M}$ into a new $\overline{M}$ which can then be applied fast.

We also explained in the previous chapter why we needed probing: to obtain an explicit approximation of D, to be compressed in order to obtain a fast matrix-vector product. Indeed, we do not have direct access to the entries of D, but rather we need to solve the costly exterior problem every time we need a multiplication of D with a vector. We have already mentioned how the approach of Lin et al. [40], for example, would require $O(\log N)$ such solves, with a large constant.

We alluded that we might be able to compress each block M (or $\widetilde{M}$) of D (or $\widetilde{D}$) when we presented background material in chapter 2. Indeed, we discussed the fact that the half-space Green's function G_{half} in constant medium is separable and low-rank away from its singularity. Because the half-space DtN map kernel K is simply two derivatives of G_{ext}, we expect K to also be separable and low-rank, and we prove this at the end of the present chapter, in section 4.3. See also the numerical verification of that theorem in section 4.4. Because the half-space DtN map is strongly related to the exterior DtN map as we mentioned in chapter 2, we expect the exterior DtN map kernel to also be separable and low rank, at least in favorable conditions such as a constant medium. But first, as in the previous chapter, we begin by explaining the technique we use, partitioned low-rank (PLR) matrices, in section 4.1. Compression of an N by N matrix into the PLR framework is nearly linear in N, and so is matrix-

vector multiplication. We then show the results of using this PLR technique on test cases in section 4.2.

4.1 Partitioned low-rank matrices

As we have discussed in chapter 2, when an operator is separable and low-rank, we expect its numerical realization to have low-rank blocks under certain conditions. In our case, the DtN map $K(x-y)$ is separable and low-rank away from the singularity $x = y$ and so we expect its numerical realization to have low-rank blocks away from its diagonal. This calls for a compression scheme such as the hierarchical matrices of Hackbush et al. [29], [27], [8], to compress off-diagonal blocks. However, because we expect higher ranks away from the singularity in variable media, and because different blocks of the DtN map will show a singularity elsewhere than on the diagonal, we decide to use a more flexible scheme called *partitioned low rank* matrices, or PLR matrices, from [34].

4.1.1 Construction of a PLR matrix

PLR matrices are constructed recursively, using a given tolerance ϵ and a given maximal rank $R_{\max}$. We start at the top level, level 0, with the matrix M which is N by N and N is a power of two[1]. We wish to compress M (in the next sections we will use this compression scheme on probed blocks $\widetilde{M}$ of the DtN map, but we use M here for notational simplicity). We first ask for the numerical rank R of M. The numerical rank is defined by the Singular Value Decomposition and the tolerance ϵ, as the number R of singular values that are larger than or equal to the tolerance. If $R > R_{\max}$, we split the matrix in four blocks and recurse to the next level, level 1 where blocks are $N/2^1$ by $N/2^1$. If the numerical rank of M is less than or equal to $R_{\max}$, $R \leq R_{\max}$, we express M in its low-rank form by truncating the SVD of M after R terms. That is, the SVD of M is $M = U\Sigma V^* = \sum_{j=1}^{N} U_j \sigma_j V_j^*$ where U and V are orthonormal matrices with columns $\{U_j\}_{j=i}^{N}$ and $\{V_j\}_{j=i}^{N}$, and Σ is the diagonal matrix of decreasing singular values: $\Sigma = \mathrm{diag}(\sigma_1, \sigma_2, \ldots, \sigma_N)$. Then, if $R \leq R_{\max}$, we compress M to $\overline{M} = \sum_{j=1}^{R} U_j \sigma_j V_j^*$ by truncating the SVD of M after R terms.

If we need to split M and recurse down to the next level, we do the following. First, we split M in four square blocks of the same size: take the first $N/2$ rows and columns to make the first block, then taking the first $N/2$ rows and last $N/2$ columns to make the second block, etc. And now we apply the step described in the previous paragraph to each block of M, checking the block's numerical rank and compressing it or splitting it depending on that numerical rank. Whenever we split up a block, we label it as "hierarchical", and call its four sub-blocks its *children*. Whenever a block was not divided, and hence compressed instead, we label it as "compressed", and we may call it a "leaf" as well.

[1] Having a square matrix with dimensions that are powers of two is not necessary, but makes the discussion easier.

If a block has dimension $R_{\max}$ by $R_{\max}$, then its numerical rank will be $R_{\max}$, and so once blocks have dimensions smaller than or equal to the maximal desired rank $R_{\max}$, we can stop recursing and store the blocks directly. However, especially if $R_{\max}$ is large, we might still be interested in compressing those blocks using the SVD. This is what we do in our code, and we label such blocks as "compressed" as well. When we wish to refer to how blocks of a certain matrix M have been divided when M was compressed in the PLR framework, or in particular to the set of all leaf blocks and their positions in M, we refer to the "structure" of M. We see then that the structure of a PLR matrix will have at most L levels, where $N/R_{\max} = 2^L$ so $L = \log_2 N/R_{\max}$.

Implementation details

Algorithm 1 presents pseudocode for the construction of a PLR matrix from a dense matrix. In practice, when we compute the SVD of a block, we use the randomized SVD[2] [39]. This allows us to use only a few matrix-vector multiplies between the block (or its transpose) and random vectors to form an approximate reduced SVD. This is a faster way of producing the SVD, and thus also of finding out whether the numerical rank of the block is larger than R. The randomized SVD requires about 10 more random matrix-vector multiplies than the desired maximal rank R. This is why, in algorithm 1, the call to *svd* has two arguments: the block we want to find the SVD of, and the maximal desired rank $R_{\max}$. The randomized SVD algorithm then uses 10 more random vectors than the quantity $R_{\max}$ and returns an SVD of rank $R_{\max} + 1$. We use the $R_{\max} + 1^{\text{st}}$ singular value in Σ to test whether we need to split the block and recurse or not.

Algorithm 1. *Compression of matrix M into Partitioned Low Rank form, with maximal rank R_{max} and tolerance ϵ*

```
1:  function H = PLR(M, R_max, ε)
2:      [U, Σ, V] = svd(M, R_max)                          ▷ Randomized SVD
3:      if ∃R ∈ {1, 2, ..., R_max} : Σ(R + 1, R + 1) < ε then
4:          Let R be the smallest such integer.
5:          H.data = {U(:, 1 : R) · Σ(1 : R, 1 : R), V(:, 1 : R)*}
6:          H.id = 'c'                                     ▷ This block is "compressed"
7:      else                                               ▷ The M_ij's are defined in the text
8:          for i = 1:2 do
9:              for j = 1:2 do
10:                 H.data{i, j} = PLR(M_ij, R_max, ε )    ▷ Recursive call
11:             end for
12:         end for
13:         H.id = 'h'                                     ▷ This block is "hierarchical"
14:     end if
15: end function
```

[2]Theoretically, this randomized SVD has a failure probability, but we can choose a parameter to make this probability on the order of 10^{-16}, and so we ignore the fact that the randomized SVD could fail.

Complexity

The complexity of this algorithm depends on the complexity of the SVD algorithm we use. The randomized SVD has complexity $O(N_B R_{\max}^2)$ where N_B is the dimension of block B whose SVD we are calculating. This can be much better, especially for larger blocks, than standard SVD algorithms which have complexity $O(N_B^3)$. The total complexity of the compression algorithm will then depend on how many blocks of which size and rank we find the randomized SVD of. We shall discuss this in more detail when we discuss also the complexity of a matrix-vector product for special structures.

Error analysis

To understand the error we make by compressing blocks in the PLR framework, we first note that compressing a block M to its R-rank approximation $\overline{M} = \sum_{j=1}^{R} U_j \sigma_j V_j^*$ by truncating its SVD, we make an error in the L_2 norm of σ_{R+1}. That is,

$$\|M - \overline{B}\|_2 = \| \sum_{j=R+1}^{N} U_j \sigma_j V_j^* \|_2 = \sigma_{R+1}.$$

Hence, by compressing a block, we make an L_2 error for that block of at most ε because we make sure that $\sigma_{R+1} \leq \varepsilon$. Of course, errors from various blocks will compound to affect the total error we make on matrix M. We shall mention this in more detail when we discuss particular structures.

The relative Frobenius error we make between M and $\overline{M}$, the compressed approximation of M, will usually be larger than ϵ because of two factors. First of all, as we just saw, the PLR compression algorithm uses the L_2 norm. To use the Frobenius norm when deciding whether to compress or divide a blcok, we would need access to all the dingular values of each block. This would be possible using typical SVD algorithm, but quite costly. Hence we use the randomized SVD, which is faster but with which we are forced to use the L_2 norm. Another factor we have yet to mention is that errors from different blocks will compound to make the total error between M and $\overline{M}$ larger than the error between any individual blocks of M and $\overline{M}$. This of course depends on how many blocks there are in the structure of nay particular PLR matrix. We will talk more about this in subsection 4.1.3, where we explore the complexity of the compression and matrix-vector algorithms of the PLR framework. But first, we introduce matrix-vector products.

4.1.2 Multiplication of a PLR matrix with a vector

To multiply a PLR matrix M with a vector v, we again use recursion. Starting at the highest level block, the whole matrix M itself, we ask whether this block has been divided into sub-blocks. If not, we multiply the block directly with the vector v. If the block has been subdivided, then we ask for each of its children whether those have been subdivided. If not, we multiply the sub-block with the appropriate

restriction of v, and add the result to the correct restriction of the output vector. If so, we recurse again. Algorithm 2 presents the pseudocode for multiplying a vector by a PLR matrix. The algorithm to left-multiply a vector by a matrix is similar, we do not show it here.

Algorithm 2. *Multiplication of a PLR matrix H with column vectors x*

```
1: function Y = MATVEC(H,x)
2:     if H.id == 'c'  then
3:         y = H.data{1}·(H.data{2}·x)
4:     else
5:         y_1 = matvec(H.data{1,1},x(1:end/2,:))
6:               +matvec(H.data{1,2},x(end/2:end,:))
7:         y_2 = matvec(H.data{2,1},x(1:end/2,:))
8:               +matvec(H.data{2,2},x(end/2:end,:))
9:         y = [ y_1 ]
               [ y_2 ]
10:    end if
11: end function
```

Complexity

The complexity of this algorithm is easily understood. Recall that we store blocks that are not divided not as a full matrix, but as column vectors corresponding to the orthonormal matrix U of the SVD, and to the product ΣV^* of the SVD. Every time we multiply such an N_B by N_B block B that has not been subdivided with the corresponding restriction of v, we first multiply the restriction $\widetilde{v}$ of v with ΣV^*, and then multiply the result with U. Let $R_B \leq R_{\max}$ be the numerical rank of that block. Then, we first make $N_B R_B$ multiplication operations and $(N_B - 1)R_B$ addition operations for the product $\Sigma V^* \widetilde{v}$, and then $R_B N_B$ multiplication operations and $(R_B - 1)N_B$ addition operations for the product of that with U. Hence we make approximately $4N_B R_B$ operations per block, where again N_B is the dimension of the block B and R_B is its numerical rank.

The total number of operations for multiplying a vector v with a PLR matrix M, then, is about

$$\sum_{B \text{ is compressed}} 4N_B R_B, \tag{4.1}$$

where we sum over all "compressed" blocks.

Evidently there is a trade-off here. Asking for a small maximal rank $R_{\max}$ may force blocks to be subdivided a lot. We then have small R_B's, and small N's, but a lot of blocks. On the other hand, having a larger $R_{\max}$ means a lot of blocks can be remain large. We then have large R_B's and N_B's, but very few blocks. We shall use the complexity count (4.1) later on to decide on which $R_{\max}$ to choose for each block of the DtN map. To have an idea of whether such a matrix-vector multiplication might result in a fast algorithm, we need to introduce new terminology regarding the structure of PLR matrices, and we do so in the next subsection.

4.1.3 Structure and complexity

As we did with matrix probing, we again split up the probed exterior DtN map $\widetilde{D}$ in submatrices corresponding to the different sides of $\partial\Omega$. We called those *blocks* of $\widetilde{D}$ in the previous chapter, but we now call them *submatrices*, to differentiate them from the blocks in the structure we obtain from compressing to a PLR matrix. So $\widetilde{D}$ is split up in submatrices $\widetilde{M}$ that represent edge-to-edge restrictions of $\widetilde{D}$. We then only have to compress once each unique submatrix $\widetilde{M}$ to obtain an approximation $\overline{M}$ of $\widetilde{M}$. We can use those compressed $\overline{M}$'s to define $\overline{D}$, an approximation of $\widetilde{D}$. If we need to multiply our final approximation $\overline{D}$ by a vector, we may then split that vector in blocks corresponding to the sides of $\partial\Omega$ and use the compressed submatrices and the required PLR matrix algebra to obtain the result with low complexity.

What is particular about approximating the DtN map on a square boundary $\partial\Omega$ is that distinct submatrices are very different. Those that correspond to the restriction of D from one edge to that same edge, and as such are on the diagonal of D, are harder to probe as we saw in the previous chapter because of the diagonal singularity. And, because of the diagonal singularity, they might be well-suited for compression by hierarchical matrices [5].

However, submatrices of D corresponding to two edges that are side by side (for example, the bottom and right edges of the boundary of $[0,1]^2$) see the effects of the diagonal of D in their upper-right or lower-left corners, and entries of such submatrices decay in norm away from that corner. Thus a hierarchical matrix would be ill-suited to compress such a submatrix. This is why the PLR framework is so useful to us: it automatically adapts to the submatrix at hand, and to whether there is a singularity in the submatrix, and where that singularity might be.

Similarly, when dealing with a submatrix of D corresponding to opposite edges of $\partial\Omega$, we see that entries with higher norm are in the upper-right and bottom-left corners, so again PLR matrices are more appropriate than hierarchical ones. However, note that because such submatrices have very small relative norm compared to D, and were probed with only one or two basis matrices in the previous chapter, their PLR structure is often trivial.

In order to help us understand the complexity of PLR compression and matrix-vector products, we first study typical structures of hierarchical [8], [30] and PLR matrices.

Weak hierarchical matrices

Definition 12. *A matrix is said to have* weak hierarchical structure *when a block is compressed if and only if its row and column indices do not overlap.*

The weak hierarchical structure of a matrix is shown in Figure 4-1. For example, let the matrix M be 8×8. Then the block at level 0 is M itself. The row indices of M are $\{1, 2, \ldots, 8\}$, and so are its column indices. Since those overlap, we divide the matrix in four. We are now at level 1, with four blocks. The $(1,1)$ block has row indices $\{1, 2, 3, 4\}$, and its column indices are the same. This block will have to be divided. The same holds for block $(2,2)$. However, block $(1,2)$ has row indices

$\{1,2,3,4\}$ and column indices $\{5,6,7,8\}$. Those two sets do not overlap, hence this block is compressed. The same will be true of block $(2,1)$.

We note that, if the matrix M has a weak hierarchical structure, we have a fast matrix-vector product. We may use the heuristic in 4.1 to obtain the complexity of that product, assuming for simplicity that all R_B's are $R_{\max}$. Hence we account for all compressed blocks, starting from the 2 larger blocks on level 1, of size $N/2$ by $N/2$ (one on each side of the diagonal): they correspond to a maximum of $4N/2 \times R_{\max}$ operations (multiplications and divisions) each, and there is two of them, so they correspond to a total of $4NR_{\max}$ operations. Then, the next larger blocks are of size $N/4$ by $N/4$, and there is 4 of them (two on each side of the diagonal). Hence they correspond to a total of $4 \times 4 \times N/4 \times R_{\max} = 4NR_{\max}$ operations. Since we have $L = \log_2 N/R_{\max}$ levels, or different block sizes, and as we can see each of those block sizes will contribute at most $4NR_{\max}$ operations, we have about $4N \log N/R_{\max} R_{\max}$ operations from off-diagonal blocks. We are left with the diagonal blocks. Those have size $R_{\max}$ and there are $N/R_{\max}$ of them, so the complexity of multiplying them by a vector is at most $4N/R_{\max} R_{\max}^2 = 4NR_{\max}$ operations. Hence the total complexity of a matrix-vector multiplication with a weak hierarchical matrix is

$$4NR_{\max} \log \frac{2N}{R_{\max}}. \tag{4.2}$$

This is clearly faster, asymptotically, than the typical complexity of a dense matrix-vector product which is of $2N^2$ operations.

As to the complexity of the compression algorithm, we do a similar calculation, but here the cost per block is $O(N_B R_{\max}^2)$, so all we need is to sum the dimensions of all blocks that we used the SVD on. We start by taking the SVD of the matrix itself, then of all the blocks on level 1, then half of the blocks on level 2, then a quarter of the blocks on level 3, etc. Hence the complexity of compression is

$$R_{\max}^2 \left(N + \sum_{l=1}^{L-1} \frac{N}{2^l} \frac{4^l}{2^{l-1}} \right) = R_{\max}^2 (N + 2N(L-1)) \leq 2NR_{\max}^2 \log \frac{N}{R_{\max}}$$

which is nearly linear in N.

Finally, we address briefly the error in the matrix M that is made when it is compressed. We especially care about the error in a matrix-vector multiplication $w = Mv$. We can see in this case that, for any entry j in w, there will be error coming from all multiplications of the appropriate restriction of v with the corresponding block intersecting row j of M. Since there are about $\log \frac{N}{R_{\max}}$ such blocks in row j, we can estimate that by giving a tolerance ε to the PLR compression algorithm, we will obtain an error in matrix-vector multiplications of about $\varepsilon \log N$. As we will see in section 4.2, dividing the "desired" error by a factor of 1 to 25 to obtain the necessary ε will work quite well for our purposes.

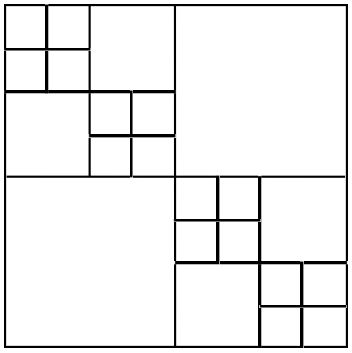

Figure 4-1: Weak hierarchical structure, $\frac{N}{R_{\max}} = 8$.

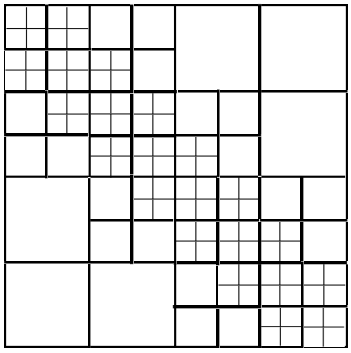

Figure 4-2: Strong hierarchical structure, $\frac{N}{R_{\max}} = 16$.

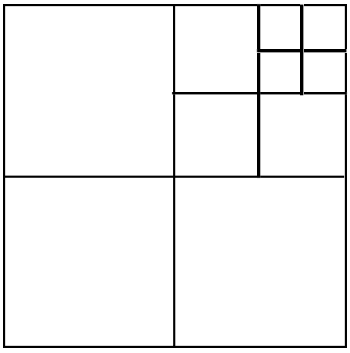

Figure 4-3: Corner PLR structure, $\frac{N}{R_{\max}} = 8$.

Strong hierarchical matrices

Next, we define a matrix with a *strong hierarchical structure*. This will be useful for matrices with a singularity on the diagonal.

Definition 13. *A matrix is said to have* strong hierarchical structure *when a block is compressed if and only if its row and column indices are separated by at least the width of the block.*

The strong hierarchical structure of a matrix is shown in Figure 4-2. We can see that, the condition for compression being stronger than in the weak case, more blocks will have to be divided. For example, let the matrix M be 8×8 again. Then the block at level 0 is M itself, and again its row and column indices overlap, so we divide the matrix in four. We are now at level 1, with four blocks. The $(1,1)$ block will still have to be divided, its row and column indices being equal. The same holds for block $(2,2)$. Now, block $(1,2)$ has row indices $\{1,2,3,4\}$ and column indices $\{5,6,7,8\}$. Those two sets do not overlap, but the distance between them, defined as the minimum of $|i-j|$ over all row indices i and column indices j for that block, is 1. Since the width of the block is 4, which is greater than 1, we have to divide the block following Definition 13. However, at level 2 which has 16 blocks of width 2, we can see that multiple blocks will be compressed: $(1,3), (1,4), (2,4), (3,1), (4,1), (4,2)$.

The matrix-vector multiplication complexity of matrices with a strong hierarchical structure is

$$12NR_{\max} \log \frac{N}{2R_{\max}}. \tag{4.3}$$

Again this is faster, asymptotically, than the typical $2N^2$ operations of a dense matrix-vector product. We can obtain this number once again by accounting for all the blocks and using 4.1. More precisely, we have $3\sum_{j=1}^{l-1} 2^l = 6\frac{1-2^{l-1}}{1-2} = 6(2^{l-1}-1)$ compressed blocks at level l, hence those blocks have size $N/2^l$. This is true for $l = 2, \dots L-1$, where again $L = \log_2 N/R_{\max}$ is the number of levels. Notice that, as expected, we do not have any compressed blocks, or leaves, at level 1. The contribution of those

blocks to the matrix-vector complexity will be

$$\begin{aligned} 4R_{\max}\sum_{l=2}^{L-1}\left(6(2^{l-1}-1)\ \frac{N}{2^l}\right) &= 12NR_{\max}\sum_{l=2}^{L-1}\frac{2^l-2}{2^l} \\ &= 12NR_{\max}(L-2-\sum_{l=0}^{L-3}\frac{1}{2}\frac{1}{2^l}) \\ &= 12NR_{\max}(L-2-(1-1/2^{L-2})) \\ &= 12NR_{\max}(L-3+1/2^{L-2}). \end{aligned}$$

We need to add to this quantity the complexity coming from the smallest blocks, of size $N/2^L$. There are

$$6(2^{L-1}-1)+2^L+2(2^L-1)=6\cdot 2^L-8$$

such blocks, and so the corresponding complexity is

$$4R_{\max}(6\cdot 2^L-8)(N/2^L)=4R_{\max}N(6-8/2^L).$$

Adding this to our previous result, we obtain the final complexity of a matrix-vector multiplication:

$$\begin{aligned} & 12NR_{\max}(L-3+1/2^{L-2})+4R_{\max}N(6-8/2^L) \\ = \ & 12NR_{\max}(L-1+\frac{1}{2^{L-2}}-\frac{2}{3\cdot 2^{L-2}}) \\ \leq \ & 12NR_{\max}L, \end{aligned}$$

as stated previously.

For the complexity of the compression algorithm, again we sum the dimensions of all blocks whose SVD we calculated: the matrix itself, the 4 blocks of level 1, the 16 blocks of level 2, 40 blocks in level 3, etc. Hence the complexity of compression is

$$\begin{aligned} R_{\max}^2\left(N+\frac{N}{2}4+\sum_{l=2}^{L}\frac{N}{2^l}(6\cdot 2^l-8)\right) &= R_{\max}^2\left(N+2N+N\sum_{l=2}^{L-1}\left(6-\frac{8}{2^l}\right)\right) \\ &= R_{\max}^2N\left(3+6(L-2)-8\ \frac{1}{4}\ \frac{1-1/2^{L-2}}{1-1/2}\right) \\ &= R_{\max}^2N\left(6L-9-4(1-1/2^{L-2})\right) \\ &= R_{\max}^2N\left(6L-13+1/2^{L-2}\right) \\ &\leq R_{\max}^2N\left(6L-12\right) \end{aligned}$$

or

$$6NR_{\max}^2\log\frac{N}{4R_{\max}}.$$

This again is nearly linear. Using similar arguments as in the weak case, we can

estimate that by giving a tolerance ε to the PLR compression algorithm, we will obtain an error in matrix-vector multiplications of about $\varepsilon \log N$ again.

Corner PLR matrices

One final structure we wish to define, now useful for matrices with a singularity in a corner, is the following:

Definition 14. *A matrix is said to have* corner PLR structure*, with reference to a specific corner of the matrix, when a block is divided if and only if both its row and column indices contain the row and column indices of the entry corresponding to that specific corner.*

Figure 4-3 shows a top-right corner PLR structure. Again, we take an 8×8 matrix M as an example. The top-right entry has row index 1 and column index 8. We see that the level 0 block, M itself, certainly contains the indices $(1, 8)$, so we divide it. On the next level, we have four blocks. Block $(1, 2)$ is the only one that has both row indices that contain the index 1, and column indices that contain the index 8, so this is the only one that is divided. Again, on level 2, we have 16 blocks of size 2, and block $(1, 4)$ is the only one divided.

As for the corner PLR matrices, their matrix-vector multiplication complexity is:

$$8NR_{\max}. \tag{4.4}$$

Indeed, we see we have 3 blocks of size $N/2^L$ for $l = 1, 2, \ldots, L-1$. This is a constant number of blocks per level, which means that matrix-vector multiplication will be even faster. We also have 4 blocks at the lowest level, of size $N/2^L$. The complexity is then

$$\begin{aligned} 4R_{\max}3\sum_{l=1}^{L-1} N/2^l \; + 4R_{\max}4N/2^L &= 4R_{\max}N(3(1-1/2^{L-1}) + 2/2^{L-1}) \\ &= 8NR_{\max}. \end{aligned}$$

For the complexity of the compression algorithm, we sum the dimensions of all blocks whose SVD we calculated: the matrix itself, the 4 blocks of level 1, 4 blocks in level 2, 4 blocks in level 3, etc. Hence the complexity of compression is

$$\begin{aligned} R_{\max}^2\left(N + 4\sum_{l=1}^{L}\frac{N}{2^l}\right) &= R_{\max}^2 N\left(1 + 4\,\frac{1}{2}\,\frac{1-1/2^L}{1-1/2}\right) \\ &= R_{\max}^2 N\left(1 + 4 - 4/2^L\right) \\ &\leq R_{\max}^2 N\,(4) \end{aligned}$$

or

$$4NR_{\max}^2.$$

Hence the complexity of compression for corner PLR matrices is linear. And again, we estimate that by giving a tolerance ε to the PLR compression algorithm, we will obtain an error in matrix-vector multiplications of about $\varepsilon \log N$.

Now that we have explained these three special structures, and how they provide a fast matrix-vector product, we are ready to discuss using PLR matrices specifically for the exterior DtN map.

4.2 Using PLR matrices for the DtN map's submatrices

As we recall, obtaining the full DtN map from solving the exterior problem $4N$ times is too costly, and so we use matrix probing to approximate the DtN map D by $\widetilde{D}$ using only a few exterior solves. If we were to try to use PLR matrices directly on D, we would have to find the SVD of many blocks. Since we do not have access to the blocks themselves, we would need to use the randomized SVD, and hence to solve the exterior problem, on random vectors restricted to the block at hand. As we mentioned before, Lin et al. have done something similar in [40], which required $O(\log N$ matrix-vector multiplies with a large constant, or in our case exterior solves. This is too costly, and this is why we use matrix probing first to obtain an approximate $\widetilde{D}$ with a small, nearly constant number of exterior solves.

Now that we have access to $\widetilde{D}$ from matrix probing, we can approximate it using PLR matrices. Indeed, we have access to the full matrix $\widetilde{D}$, and so finding the SVD of a block is not a problem. In fact, we use the randomized SVD for speed, not because we only have access to matrix-vector multiplies.

Compressing one of those edge-to-edge submatrices under the PLR matrix framework requires that we pick both a tolerance ϵ and a maximal desired rank $R_{\max}$. We explain in the next subsections how to choose appropriate values for those parameters.

4.2.1 Choosing the tolerance

Because our submatrices come from probing, they already have some error attached to them, that is, the relative probing error as defined in equation 3.8 of chapter 3. Therefore, it would be wasteful to ask for the PLR approximation to do any better than that probing error.

Also, when we compress blocks in the PLR compression algorithm, we make an absolute error in the L_2 norm. However, because of the high norm of the DtN map, it makes more sense to consider the relative error. We can thus multiply the relative probing error we made on each submatrix $\widetilde{M}$ by the norm of the DtN map D to know the absolute error we need to ask of the PLR compression algorithm. And since the L_2 norm is smaller than the Frobenius norm, and errors from each block compound, we have found empirically that asking for a tolerance ϵ which is a factor of 1 to 1/100 of the absolute probing error of a submatrix works for obtaining a similar Frobenius error from the PLR approximation. As a rule of thumb, this factor needs

to be smaller for diagonal submatrices M of D, but can be equal to 1 for submatrices corresponding to opposite edges of $\partial\Omega$.

Of course, we do not want to use an ε which is too smaller either. That might force the PLR compression algorithm to divide blocks more than needed, and make the matrix-vector multiplications slower than needed.

4.2.2 Minimizing the matrix-vector application time

Our main objective in this chapter is to obtain a fast algorithm. To this end, we try to compress probed submatrices of the DtN map using various values of the parameter $R_{\max}$, and choose the value that will give us the fastest matrix-vector multiplies. We use the known complexity of a matrix-vector multiplication (4.1) to find the rank $R_{\max}$ that minimizes the complexity, from doing a few tests, and we use the compressed submatrix corresponding to that particular maximal rank in our Helmholtz solver. A different choice of complexity might be used depending on the operating system and coding language used, since slow downs might occur because of cache size, operations related to memory, matrix and vector operations, etc.

However, we note that we may compare the actual complexity from the particular structure obtained by PLR compression to the "ideal" complexities coming from the special structures we have mentioned before. Indeed, for a submatrix on the diagonal, we can compare its matrix-vector complexity to that of weak and strong hierarchical matrices. That will give us an idea of whether we have a fast algorithm. One thing we notice in most of our experiments is that, for diagonal blocks, the actual complexity usually becomes smaller as $R_{\max}$ increases, until we arrive at a minimum with the $R_{\max}$ that gives us the better compromise between too many and blocks of too high a rank. Then, the complexity increases again. However, the complexity increases slower than that of both weak and strong matrices.

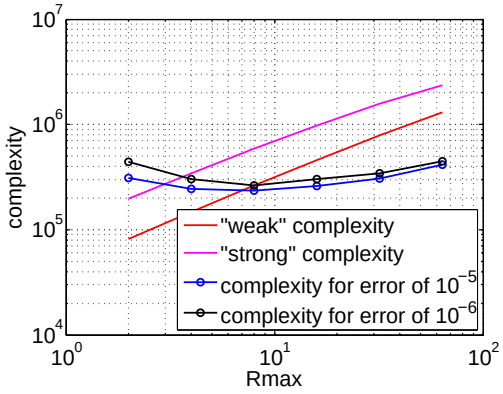

Figure 4-4: Matrix-vector complexity for submatrix $(1,1)$ for $c \equiv 1$, various $R_{\max}$. Probing errors of 10^{-5}, 10^{-6}.

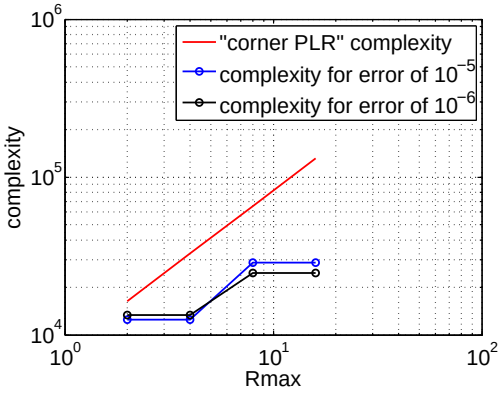

Figure 4-5: Matrix-vector complexity for submatrix $(2,1)$ for $c \equiv 1$, various $R_{\max}$. Probing errors of 10^{-5}, 10^{-6}.

Figure 4-4 confirms this phenomenon for the $(1,1)$ block of the constant medium. From this figure, we would then pick $R_{\max} = 8$ since this is the value of $R_{\max}$ that

corresponds to the smallest actual complexity of a matrix-vector product, both for a relative probing error of 10^{-5} and 10^{-6}. Figure 4-5 confirms this phenomenon as well for the $(2,1)$ block of the constant medium. From this figure, we would then pick $R_{\max} = 4$ again both for a relative probing error of 10^{-5} and 10^{-6}, since this is the value of $R_{\max}$ that corresponds to the smallest actual complexity of a matrix-vector product (it is hard to tell from the figure, but the complexity for $R_{\max} = 2$ is just larger than that for $R_{\max} = 4$ in both cases).

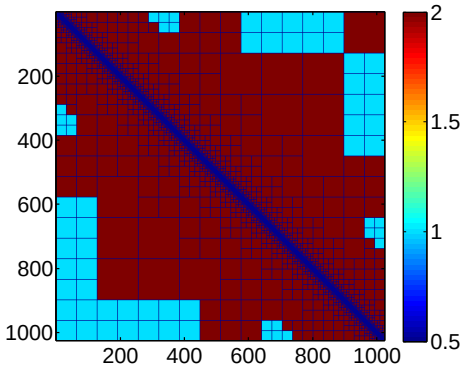

Figure 4-6: PLR structure of the probed submatrix $(1,1)$ for $c \equiv 1$, $R_{\max} = 2$. Each block is colored by its numerical rank.

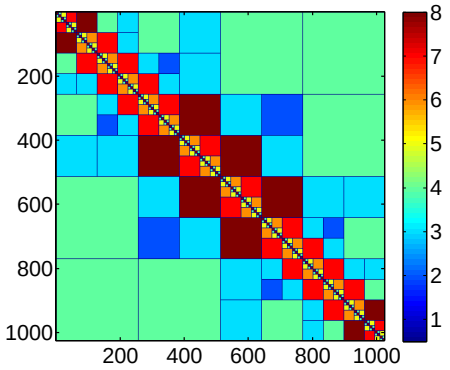

Figure 4-7: PLR structure of the probed submatrix $(1,1)$ for $c \equiv 1$, $R_{\max} = 8$. Each block is colored by its numerical rank.

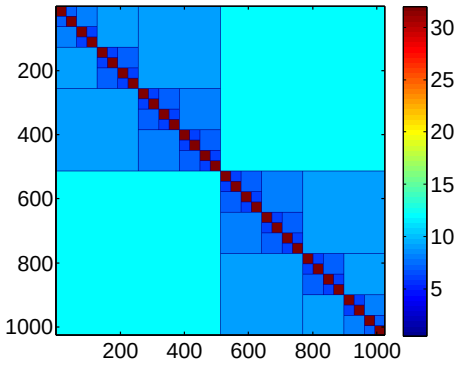

Figure 4-8: PLR structure of the probed submatrix $(1,1)$ for $c \equiv 1$, $R_{\max} = 32$. Each block is colored by its numerical rank.

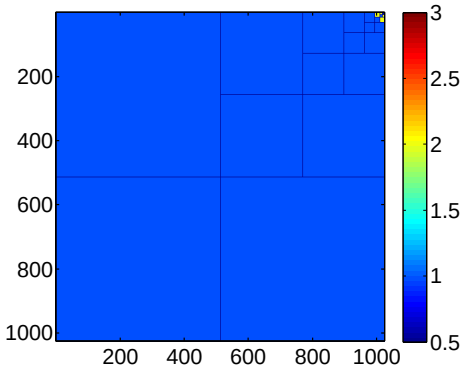

Figure 4-9: PLR structure of the probed submatrix $(2,1)$ for $c \equiv 1$, $R_{\max} = 4$. Each block is colored by its numerical rank.

It is also informative to have a look at the structures we obtain, and the ranks of each block. Figures 4-6, 4-7 and 4-8 refer to block $(1,1)$ again of $c \equiv 1$, for different values of $R_{\max}$. As we expect, having $R_{\max} = 2$ in Figure 4-6 in this case forces blocks to be very small, which is wasteful. On the other hand, a larger $R_{\max} = 32$ in Figure

4-8 is not better because then we have fewer blocks, but they have large rank. Still, because the structure is that of a weak hierarchical matrix, with blocks that have in fact rank smaller than $R_{\max}$, we obtain a fast matrix-vector product. However, the ideal is $R_{\max} = 8$ in Figure 4-7, which minimizes the complexity of a matrix-vector product by finding the correct balance between fewer blocks but small rank. We see it almost has a strong hierarchical structure, but in fact with more large blocks and fewer small blocks. As for the $(2, 1)$ block, we see its PLR structure in Figure 4-9: it is actually a corner PLR structure but the numerical ranks of blocks are always lower than $R_{\max} = 4$ so the matrix-vector multiplication of that submatrix will be even faster than for a corner PLR structure, as we knew from Figure 4-5.

4.2.3 Numerical results

Table 4.1: PLR compression results, $c \equiv 1$

$R_{\max}$ for $(1,1)$	$R_{\max}$ for $(2,1)$	$\|D - \overline{D}\|_F / \|D\|_F$	$\|u - \overline{u}\|_F / \|u\|_F$	Speed-up
2	2	$4.2126e-01$	$6.5938e-01$	115
2	2	$4.2004e-02$	$7.3655e-02$	93
2	2	$1.2517e-03$	$2.4232e-03$	55
4	2	$1.1210e-04$	$4.0003e-04$	42
8	4	$1.0794e-05$	$1.4305e-05$	32
8	4	$6.5496e-07$	$2.1741e-06$	29

Table 4.2: PLR compression results, c is the Gaussian waveguide.

$R_{\max}$ for $(1,1)$	$R_{\max}$ for $(2,2)$	$R_{\max}$ for $(2,1)$	$\|D - \overline{D}\|_F / \|D\|_F$	$\|u - \overline{u}\|_F / \|u\|_F$	Speed-up
2	2	2	$6.6034e-02$	$1.4449e-01$	105
2	2	2	$1.8292e-02$	$7.4342e-02$	74
2	2	2	$2.0948e-03$	$1.1014e-02$	59
4	4	2	$2.3740e-04$	$1.6023e-03$	47
8	4	4	$1.5369e-05$	$8.4841e-05$	36
8	8	4	$3.4148e-06$	$1.7788e-05$	30

We have compressed probed DtN maps and used them in a Helmholtz solver with success. We have used the same probed matrices as in the previous chapter, and so we refer the reader to Tables 3.1, 3.2, 3.3, 3.4, 3.5, 3.6, 3.7 for all the parameters we used then.

We now present results for PLR compression in a Helmholtz solver in Tables 4.1, 4.2, 4.3, 4.4, 4.5, 4.6. For each medium, we show the chosen $R_{\max}$ of the most important (in norm) submatrices. For all other submatrices, $R_{\max} \leq 2$. We then show the relative norm of the error between the PLR compression $\overline{D}$ and the actual DtN map D. We also show the relative error between the solution $\overline{u}$ computed using $\overline{D}$

Table 4.3: PLR compression results, c is the Gaussian slow disk.

$R_{\max}$ for $(1,1)$	$R_{\max}$ for $(2,1)$	$\|D-\overline{D}\|_F/\|D\|_F$	$\|u-\overline{u}\|_F/\|u\|_F$	Speed-up
2	2	$9.2307e-02$	$1.2296e+00$	97
2	2	$8.1442e-03$	$4.7922e-02$	69
4	2	$1.2981e-03$	$3.3540e-02$	44
4	2	$1.1680e-04$	$1.0879e-03$	39
4	2	$2.5651e-05$	$1.4303e-04$	37

Table 4.4: PLR compression results, c is the vertical fault, sources on the left and on the right.

$R_{\max}$ for $(1,1)$	$R_{\max}$ for $(2,2)$	$\frac{\|D-\overline{D}\|_F}{\|D\|_F}$	$\frac{\|u-\overline{u}\|_F}{\|u\|_F}$, left	$\frac{\|u-\overline{u}\|_F}{\|u\|_F}$, right	Speed-up
2	2	$2.6972e-01$	$5.8907e-01$	$4.6217e-01$	105
2	2	$9.0861e-03$	$3.9888e-02$	$2.5051e-02$	67
1	4	$8.7171e-04$	$3.4377e-03$	$2.4279e-03$	53

in the Helmholtz solver and the actual solution u using D as the DtN map. Finally, we show the "speed-up" obtained from taking the ratio of the complexity of using a dense matrix-vector product for $\widetilde{D}$, which would be of about $2 \times 16N^2$, to the total complexity of a matrix-vector product of $\bar{D}$. This ratio tells us how much faster than a dense product the PLR compression is. We see that this ratio ranges from about 50 to 100 for all media expect the periodic media, with a smaller ratio associated with asking for a higher accuracy, as expected. The speed-up ratio is between 10 and 30 for the periodic media, but as we recall the value of N here is smaller: $N = 320$. Larger values of N should lead to a better speed-up.

4.3 The half-space DtN map is separable and low rank: theorem

As we have mentioned before, the Green's function for the half-space Helmholtz equation is separable and low rank [20]. We investigate here the half-space DtN map kernel, which is related to the Green's function through two derivatives, as we saw in section 2.3.3, and we obtain a similar result to that of [20]. We state the result here, and prove it in the next subsections. We then end this section with a discussion on generalizing our theorem for heterogeneous media.

Let $\mathbf{x} = (x, 0)$ and $\mathbf{y} = (y, 0)$ be points along the half-space boundary, $x \neq y$. Recall the Dirichlet-to-Neumann map kernel for the half-space Helmholtz equation (2.9) with homogeneous medium $c \equiv 1$ and $\omega = k/c$ is:

$$K(|x-y|) = \frac{ik^2}{2} \frac{H_1^{(1)}(k|x-y|)}{k|x-y|}. \tag{4.5}$$

Table 4.5: PLR compression results, c is the diagonal fault.

$R_{\max}$ for $(1,1)$	$R_{\max}$ for $(2,2)$	$\|D-\overline{D}\|_F/\|D\|_F$	$\|u-\overline{u}\|_F/\|u\|_F$	Speed-up
2	2	$1.4281e-01$	$5.3553e-01$	98
2	2	$1.9108e-02$	$7.8969e-02$	76
2	4	$2.5602e-03$	$8.7235e-03$	49

Table 4.6: PLR compression results, c is the periodic medium.

$R_{\max}$ for $(1,1)$	$R_{\max}$ for $(2,1)$	$\|D-\overline{D}\|_F/\|D\|_F$	$\|u-\overline{u}\|_F/\|u\|_F$	Speed-up
2	2	$1.2967e-01$	$2.1162e-01$	32
2	2	$3.0606e-02$	$5.9562e-02$	22
8	2	$9.0682e-03$	$2.6485e-02$	11

Theorem 3. *Let $0 < \epsilon \leq 1/2$, and $0 < r_0 < 1$, $r_0 = \Theta(1/k)$. There exists an integer J, functions $\{\Phi_j, \chi_j\}_{j=1}^{J}$ and a number C such that we can approximate $K(|x-y|)$ for $r_0 \leq |x-y| \leq 1$ with a short sum of smooth separated functions:*

$$K(|x-y|) = \sum_{j=1}^{J} \Phi_j(x)\chi_j(y) + E(x,y) \tag{4.6}$$

where $|E(x,y)| \leq \epsilon$, and $J \leq C\,(\log k \max(|\log \epsilon|, \log k))^2$ with C which does not depend on k, or ϵ. C does depend weakly on r_0 through the constant quantity $r_0 k$, for more details see Remark 4.3.6.

To prove this, we shall first consider the Hankel function $H_1^{(1)}$ in (4.5), and see that it is separable and low rank. Then, we shall look at the factor of $1/k|x-y|$ in (4.5), and see the need for using quadratures on a dyadic partition of the interval $[r_0, 1]$ in order to prove that this factor is also separable and low rank. Finally, we prove Theorem 3 and make a few remarks.

4.3.1 Treating the Hankel function

From Lemmas 7 and 8 of [42] we know that $H_1^{(1)}(k|x-y|)$ is separable and low-rank as a function of x and y. Looking in particular at Lemma 7 from [42], we make a slight modification of the proof to obtain the following lemma.

Lemma 2. *Slight modification of Lemmas 7 and 8, [42]. Let $0 < \epsilon \leq 1/2$, $k > 0$ and $r_0 > 0$, $r_0 = \Theta(1/k)$. Let $|x-y| > r_0$. Then there exists an integer J_1, a number C_1, and functions $\left\{\Phi_j^{(1)}, \chi_j^{(1)}\right\}_{j=1}^{J_1}$ such that*

$$H_1^{(1)}(k|x-y|) = \sum_{j=1}^{J_1} \Phi_j^{(1)}(x)\chi_j^{(1)}(y) + E_{J_1}^{(1)}(x,y), \tag{4.7}$$

where $|E^{(1)}_{J_1}(x,y)| \leq \epsilon$*, and* $J_1 \leq C_1 \log k |\log \epsilon|$ *with* C_1 *which does not depend on* k*, or* ϵ*.* C_1 *does depend weakly on* r_0 *through the quantity* $r_0 k$*. Again, see Remark 4.3.6.*

4.3.2 Treating the $1/kr$ factor

We now show that the $1/kr$ factor is also separable and low rank[3]. Notice:

$$\int_0^\infty e^{-krt} dt = \left. \frac{e^{-krt}}{-kr} \right|_0^\infty = 0 + \frac{e^{-kr0}}{kr} = \frac{1}{kr} \tag{4.8}$$

and

$$\int_0^\infty e^{-krt} dt = \int_0^T e^{-krt} dt + \int_T^\infty e^{-krt} dt, \tag{4.9}$$

where

$$\int_T^\infty e^{-krt} dt = \frac{e^{-krT}}{kr}.$$

Equation (4.8) means we can write $1/kr$ as an integral, and equation (4.9) means we can split this integral in two, one part being a definite integral, the other indefinite. But we can choose T so that the indefinite part is smaller than our error tolerance:

$$\left| \int_T^\infty e^{-krt} dt \right| \leq \epsilon.$$

For this, we consider that $\frac{e^{-krT}}{kr} \leq \frac{e^{-krT}}{C_0} \leq \epsilon$ and so we need $krT \geq |\log C_0| + |\log \epsilon|$ or $T \geq (|\log C_0| + |\log \epsilon|)/C_0$ or

$$T = O(|\log \epsilon|). \tag{4.10}$$

If we assume (4.10) holds, then we can use a Gaussian quadrature to obtain a low-rank, separable expansion of $1/kr$:

$$\frac{1}{kr} \approx \int_0^T e^{-krt} dt \approx \sum_{p=1}^n w_p e^{-krt_p}$$

where the w_p are the Gaussian quadrature weights and the t_p are the quadrature points. To determine n, the number of quadrature weights and points we need for an accuracy of order ϵ, we can use the following Gaussian quadrature error estimate [48]

[3] A different technique than the one presented here would be to expand $1/|x-y|$ using the Taylor expansion for $x > y > 0$:

$$\frac{1}{x} \frac{1}{1 - y/x} = \frac{1}{x} \left(1 + \frac{y}{x} + \left(\frac{y}{x} \right)^2 + \dots \right).$$

However, making an error of ε requires $k \log \varepsilon$ terms in the expansion because the error is large when $y/x \approx 1$ or $y \approx x$ or $r \approx r_0$.

on the interval $[a, b]$:

$$\frac{(b-a)^{2n+1}(n!)^4}{(2n+1)[(2n)!]^3} f^{(2n)}(\xi) \tag{4.11}$$

where f is the integrand, and ξ is within the bounds of integration: $f(\xi) = e^{-kr\xi}$ and $a \le \xi \le b$ where here $a = 0$ and $b = T$. Clearly

$$f^{(2n)}(\xi) = (-kr)^{2n} e^{-kr\xi}.$$

The worst case will be when $\xi = 0$ and $r = 1$ so

$$\max_{0 \le \xi \le T} \left| f^{(2n)}(\xi) \right| = (k)^{2n}.$$

We can put this back in the error estimate (4.11), using Stirling's approximation [1]

$$\sqrt{2\pi n}\ n^n e^{-n} \le n! \le e\sqrt{n}\ n^n e^{-n},$$

for the factorials, to get:

$$\begin{aligned}
\left| \frac{(b-a)^{2n+1}(n!)^4}{(2n+1)[(2n)!]^3} f^{(2n)}(\xi) \right| &\le \frac{T^{2n+1}(n!)^4}{(2n+1)[(2n)!]^3}(k)^{2n} \\
&\le \frac{T^{2n+1} e^4 (n)^2 (n)^{4n} e^{6n}}{(2n+1) e^{4n} (2\pi 2n)^{3/2} (2n)^{6n}}(k)^{2n} \\
&\le \frac{T^{2n+1} e^4 (n)^{1/2} (ke)^{2n}}{(2n+1)\pi^{3/2} (2)^{6n+3} n^{2n}} \\
&\le \frac{Te^4}{16\sqrt{n}\pi^{3/2}} \left(\frac{Tke}{8n} \right)^{2n}.
\end{aligned}$$

This is problematic because in order for the quadrature scheme to converge, we are forced to have $n > Tke/8 \approx k \log \epsilon$, which is prohibitively large. This can be understood as the difficulty of integrating accurately a function which has large higher derivatives such as this sharp exponential over a large domain such as the interval $[0, \log \epsilon]$. To solve this problem, we make a dyadic partition of the $[0, T]$ interval in $O(\log k)$ subintervals, each of which will require $O(\log \epsilon)$ quadrature points. Before we get to the details, let us redo the above error analysis for a dyadic interval $[a, 2a]$. The maximum of $\left| f^{(2n)}(\xi) \right| = (kr)^{2n} e^{-kr\xi}$ as a function of kr occurs when $kr = 2n/\xi$, and the maximum of that as a function of ξ is when $\xi = a$, so the maximum is $\left| f^{(2n)}(a) \right| = (2n/a)^{2n} e^{-(2n/a)*a}$. We can put this back in the error estimate (4.11) to

get

$$\begin{aligned}\left|\frac{(2a-a)^{2n+1}(n!)^4}{(2n+1)[(2n)!]^3}f^{(2n)}(\xi)\right| &\leq \frac{a^{2n+1}(n!)^4}{(2n+1)[(2n)!]^3}(2n/a)^{2n}e^{-2n} \\ &\leq \frac{ae^4(n)^2(n)^{4n}e^{6n}}{(2n+1)e^{4n}(2\pi 2n)^{3/2}(2n)^{6n}}(2n)^{2n}e^{-2n} \\ &\leq \frac{1.22a(n)^{1/2}}{(2n+1)(2)^{4n}} \\ &\leq \frac{a}{\sqrt{n}(2)^{4n}}.\end{aligned}$$

To have this error less than ϵ, we thus need

$$4n\log 2 \geq |\log\epsilon| + \log a/\sqrt{n}, \tag{4.12}$$

with $a \leq T \approx |\log\epsilon|$, and we see that in fact

$$n = |\log\epsilon| \tag{4.13}$$

will work.

Remark 4.3.3. *We found the maximum of* $\left|f^{(2n)}(\xi)\right| = (kr)^{2n}e^{-kr\xi}$ *to be when both* $kr = 2n/\xi$ *and* $\xi = a$*. However, we neek* $kr \leq k$*, so that we need* $a \geq 2n/k = 2|\log\varepsilon|/k$*. In the next subsection, we make sure that* $a \geq 2|\log\varepsilon|/k$ *by having the lower of the interval* I_1 *equal to* $2|\log\varepsilon|/k$*.*

4.3.4 Dyadic interval for the Gaussian quadrature

Now we are ready get into the details of how we partition the interval. The subintervals are:

$$\begin{aligned}I_0 &= \left[0, \frac{2|\log\epsilon|}{k}\right] \\ I_j &= \left[\frac{2^j|\log\epsilon|}{k}, \frac{2^{j+1}|\log\epsilon|}{k}\right], \qquad j = 1, \ldots, M-1\end{aligned}$$

where $T = \frac{2^M|\log\epsilon|}{k} = O(|\log\epsilon|)$ which implies that

$$M = O(\log k). \tag{4.14}$$

Then, for each interval I_j with $j \geq 1$, we apply a Gaussian quadrature as explained above, and we need $n = |\log\epsilon|$ quadrature points to satisfy the error tolerance of ϵ. As for interval I_0, we return to the Gaussian quadrature error analysis, where this time again k^{2n} is the maximum of $\left|f^{(2n)}(\xi)\right|$ for $\xi \in I_0$. Thus we have that the

quadrature error is:

$$\begin{aligned}\left|\frac{(2|\log\epsilon|/k-0)^{2n+1}(n!)^4}{(2n+1)[(2n)!]^3}f^{(2n)}(\xi)\right| &\leq \frac{(2|\log\epsilon|/k)^{2n+1}(n!)^4}{(2n+1)[(2n)!]^3}k^{2n}\\ &\leq \frac{(2|\log\epsilon|/k)^{2n+1}e^4(n)^2(n)^{4n}e^{6n}}{(2n+1)e^{4n}(2\pi 2n)^{3/2}(2n)^{6n}}k^{2n}\\ &\leq \frac{2^{2n+1}|\log\epsilon|^{2n+1}e^{2n}e^4n^{1/2}}{k(2n+1)(2)^{6n+3}n^{2n}}\\ &\leq \frac{2|\log\epsilon|e^4\sqrt{n}}{8k(2n+1)}\left(\frac{2|\log\epsilon|e}{8n}\right)^{2n}\end{aligned}$$

and $n = O(|\log\epsilon|)$ will satisfy the error tolerance.

To recap, we have approximated the function $1/kr$, as a function of r, by a low-rank separable expansion with error ϵ:

$$\frac{1}{k|x-y|} = \sum_{j=1}^{J_2} w_j e^{-k|x-y|t_j} + E_{J_2}^{(2)}(x,y),$$

where $J_2 = O(\log k|\log\epsilon|)$ (again, from using $O(\log k)$ intervals with $O(|\log\epsilon|)$ quadrature points on each interval), $C_0 \leq k|x-y| \leq k$, and $|E_{J_2}^{(2)}(x,y)| < \epsilon$. Clearly this expansion is separable: depending on the sign of $(x-y)$, we have $e^{-k|x-y|t_j} = e^{-kxt_j}e^{kyt_j}$ or $e^{-k|x-y|t_j} = e^{kxt_j}e^{-kyt_j}$. Either way, the exponential has been expressed as a product of a function of x only and a function of y only. Thus we have the following lemma.

Lemma 3. *Let $0 < \epsilon$, $k > 0$ and $r_0 > 0$, $r_0 = \Theta(1/k)$. Let $|x-y| > r_0$. Then there exists an integer J_2, a number C_2, and functions $\left\{\Phi_j^{(2)}, \chi_j^{(2)}\right\}_{j=1}^{J_2}$ such that*

$$\frac{1}{k|x-y|} = \sum_{j=1}^{J_2} \Phi_j^{(2)}(x)\chi_j^{(2)}(y) + E_{J_2}^{(2)}(x,y), \tag{4.15}$$

where $|E_{J_2}^{(2)}(x,y)| \leq \epsilon$, and $J_2 \leq C_2 \log k|\log\epsilon|$ with C_2 which does not depend on k, or ϵ. C_2 does depend weakly on r_0 through the constant quantity $r_0 k$. Again, see Remark 4.3.6.

4.3.5 Finalizing the proof

We now come back to the DtN map kernel K in (4.5). Using Lemmas 2 and 3, we can write each factor of K in its separable expansion:

$$\begin{aligned}
\frac{K(|x-y|)}{ik^2/2} &= H_1^{(1)}(k|x-y|)\ \frac{1}{k|x-y|} \\
&= \left(\sum_{j=1}^{J_1}\Phi_j^{(1)}(x)\chi_j^{(1)}(y) + E_{J_1}^{(1)}(x,y)\right)\left(\sum_{j=1}^{J_2}\Phi_j^{(2)}(x)\chi_j^{(2)}(y) + E_{J_2}^{(2)}(x,y)\right) \\
&= \frac{K_{(J_1,J_2)}(|x-y|)}{ik^2/2} + E_{J_2}^{(2)}(x,y)\sum_{j=1}^{J_1}\Phi_j^{(1)}(x)\chi_j^{(1)}(y) + E_{J_1}^{(1)}(x,y)\sum_{j=1}^{J_2}\Phi_j^{(2)}(x)\chi_j^{(2)}(y) \\
&+ E_{J_1}^{(1)}(x,y)E_{J_2}^{(2)}(x,y)
\end{aligned}$$

where

$$\begin{aligned}
K_{(J_1,J_2)}(|x-y|) &= \frac{ik^2}{2}\left(\sum_{j=1}^{J_1}\Phi_j^{(1)}(x)\chi_j^{(1)}(y)\right)\left(\sum_{j=1}^{J_2}\Phi_j^{(2)}(x)\chi_j^{(2)}(y)\right) \qquad (4.16) \\
&= \frac{ik^2}{2}\sum_{j=1}^{J_1J_2}\Phi_j(x)\chi_j(y). \qquad (4.17)
\end{aligned}$$

It follows that

$$\left|\frac{K-K_{(J_1,J_2)}}{ik^2/2}\right| \le \left|E_{J_2}^{(2)}\right|\left|\sum_{j=1}^{J_1}\Phi_j^{(1)}(x)\chi_j^{(1)}(y)\right| + \left|E_{J_1}^{(1)}\right|\left|\sum_{j=1}^{J_2}\Phi_j^{(2)}(x)\chi_j^{(2)}(y)\right| + \left|E_{J_1}^{(1)}\right|\left|E_{J_2}^{(2)}\right| \qquad (4.18)$$

Now clearly

$$\max_{C_0\le kr\le k}\frac{1}{kr} = \frac{1}{C_0}.$$

We also have from Lemma 3 of [16] that

$$\left|e^{-ikr}H_1^{(1)}(kr)\right| \le C(kr)^{-1/2}, \qquad kr \ge C_0$$

for some constant C which does not depend on k. Then, we have that

$$\max_{kr\ge C_0}\left|H_1^{(1)}(kr)\right| \le \frac{C}{C_0^{1/2}}.$$

What we have shown is that the quantities $1/kr$ and $H_1^{(1)}(kr)$ are bounded by some constant, call it $\widetilde{C}$, for the range of kr we are interested, that is, $C_0 \le kr \le k$. We can now go back to (4.18), using our approximations from Lemmas 2 and 3 which make an absolute error of no more than ϵ, and see that

$$\left|\frac{K-K_{(J_1,J_2)}}{ik^2/2}\right| \le 2\epsilon(\widetilde{C}+\epsilon)+\epsilon^2$$

or

$$\left|\frac{K - K_{(J_1,J_2)}}{ik^2/2}\right| = O(\epsilon). \tag{4.19}$$

Note that the expansion of $K_{(J_1,J_2)}$ in (4.16) now contains $J_1 J_2 = O\left((\log k|\log\epsilon|)^2\right)$ terms. In order to obtain an absolute error bound, we need to multiply through with $ik^2/2$ in (4.19). With $\epsilon \to \epsilon k^2$, we have now finally shown that

$$\left|K - K_{(J_1,J_2)}\right| \leq \epsilon$$

with the expansion of $K_{(J_1,J_2)}$ in (4.16) containing

$$J = J_1 J_2 = O\left((\log k(|\log\epsilon| + 2\log k))^2\right)$$

terms. We thus conclude that the DtN map is low-rank and separable away from the diagonal, with the prescriptions of Theorem 3:

$$K(|x-y|) = \frac{ik^2}{2}\frac{H_1^{(1)}(k|x-y|)}{k|x-y|} = \sum_{j=1}^{J} \Phi_j(x)\chi_j(y) + E(x,y) \tag{4.20}$$

where $|E(r)| \leq \epsilon$ for $C_0 \leq kr \leq k$ and there is a number C which does not depend on k or ϵ such that $J \leq C(\log k \max(|\log\epsilon|, \log k))^2$.

Remark 4.3.6. *We can understand the number J of ranks as made of two terms, one which is $(C_1 \log k)(C_2|\log\epsilon|)$, the other $(C_1 \log k)(\log C_3 k)$. The numbers C_1 and C_3, but not C_2, also weakly depend on the separation C_0, in the sense that the larger the separation is, the smaller those numbers are. First, we note from the discussion before equation* (4.10) *that T is smaller when C_0 is bigger. Then, a smaller C_1 comes from the discussion before equation* (4.14)*. The fact that C_2 does* not *depend on C_0 comes from the discussion after equation* (4.12)*. As for C_3, we can understand its* very weak *dependence on C_0 by looking at equation* (4.12) *and plugging in $a = T$, remembering how T depends on C_0. Thus both terms in J should behave somewhat similarly as C_0 changes. Physically, we know a greater separation means we are farther away from the singularity of the DtN map, and so we expect the map to be smoother there, and hence have lower rank.*

Remark 4.3.7. *In our numerical verifications, we have not noticed the square power in J. Rather, we observe in general that $J \sim \log k|\log\epsilon|$. The only exceptions to this behavior that we observed were for larger ϵ, such as $\epsilon = 1/10$ or sometimes $1/100$, where $J \sim \log k \log k$. From Remark 4.3.6, we know J is made up of two terms, and it makes sense that the term $\sim \log k \log k$ might become larger than the term $\sim \log k|\log\varepsilon|$ when ε is large.*

Remark 4.3.8. *We also note that in our numerical verifications, we use r_0 as small as h, which is smaller than the $r_0 \sim 1/k \sim h^{2/3}$ we prove the theorem with. If we used $r_0 \sim h$ in the theorem, this would mean $C_0 \sim n^{-1/3} < 1$. By Remark 4.3.6, this would affect J: the power of the $|\log\epsilon|$ factor would go from 2 to 3. Again, we do not notice such a higher power in the numerical simulations.*

4.3.9 The numerical low-rank property of the DtN map kernel for heterogeneous media

We would like to know as well if we can expect the half-space DtN map kernel to be numerically low-rank in heterogeneous media. We saw in section 3.4 how the half-space DtN map, in constant medium, consists of the amplitude $H(r)$, singular at $r = 0$, multiplied by the complex exponential e^{ikr}. Because of the geometrical optics expansion (3.5) of the Green's function G for the Helmholtz equation free-space problem in heterogeneous media, we expect G to have an amplitude $A(\mathbf{x}, \mathbf{y})$, which is singular at $\mathbf{x} = \mathbf{y}$, multiplied by a complex exponential with a phase corresponding to the traveltime between points $\mathbf{x}$ and $\mathbf{y}$: $e^{i\omega\tau(\mathbf{x},\mathbf{y})}$. We can expect to be able to treat the amplitude in the same way as we did before, and approximate it away from the singularity with a low-rank separable expansion. However, the complex exponential is harder to analyze because of the phase $\tau(\mathbf{x}, \mathbf{y})$, which is not so simple as $|\mathbf{x} - \mathbf{y}|$.

However, a result of [9] still allows us to find a separable low-rank approximation of a function such as $e^{i\omega\tau(\mathbf{x},\mathbf{y})}$. We refer the reader to Theorem 3.1 of [9] for the details of the proof, and simply note here the main result. Let X, Y be bounded subsets of $\mathbf{R}^2$ such that we only consider $\mathbf{x} \in X$ and $\mathbf{y} \in Y$. The width of X (or Y) is defined as the maximal Euclidean distance between any two points in that set. Then Theorem 3.1 of [9] cites that $e^{i\omega\tau(\mathbf{x},\mathbf{y})}$, $\mathbf{x} \in X$ and $\mathbf{y} \in Y$, is numerically low-rank with rank $O(|\log^4 \varepsilon|)$, given that the product of the widths of X and Y is less than $1/k$.

This translates into a restriction on how large can the off-diagonal blocks of the matrix D be while still being low-rank. Since we use square blocks in the PLR compression algorithm, we expect blocks to have to remain smaller than $1/\sqrt{k}$, equivalent to $N/\sqrt{k}$ points, in variable media. If $N^{2/3} \sim k$ as we have in this thesis because of the pollution effect, this translates into a maximum expected number of blocks of $N^{1/3}$. If we kept instead $N \sim k$, then the maximal number of blocks would be $\sqrt{N}$.

This is why, again, using PLR matrices for compressing the DtN map makes more sense than using hierarchical matrices: the added flexibility means blocks will be divided only where needed, in other words only where the traveltime τ requires blocks to be smaller in order to have low ranks. And as we saw in section 4.2, where we presented our numerical results, ranks indeed remain very small, between 2 and 8, for off-diagonal blocks of the submatrices of the exterior DtN map, even in heterogeneous media.

4.4 The half-space DtN map is separable and low rank: numerical verification

We first compute the half-space DtN map for various $k \sim N^{2/3}$, which ensures a constant error from the finite difference discretization (FD error) as we saw in section 2.2.3. We also choose a pPML width consistent with the FD error level. Then, we compute the maximum off-diagonal ranks for various fixed separations from the diagonal, that is, various r_0 such that $r \geq r_0$. To compute the ranks of a block, we fix a tolerance ϵ, find the Singular Value Decomposition of that block, and discard

all singular values smaller than that tolerance. The number of remaining singular values is the numerical rank of that block with tolerance ϵ (the error we make in Frobenius norm is not larger than ϵ). Then, the maximum off-diagonal rank for a given separation r_0 is the maximum rank of any block whose entries correspond to $r \geq r_0$. Hence we consider all blocks that have $|i-j| \geq r_0/h$, or $i-j \geq r_0/h$ with $i > j$ since the DtN map is symmetric (and so is its numerical realization, up to machine precision).

4.4.1 Slow disk and vertical fault

The Figures 4-10, 4-11, 4-12, 4-13 show the relationship between the ranks and N or ε for the slow disk, FD error of 10^{-3} and separations of $r_0 = 1$ and $r_0 = 4$.

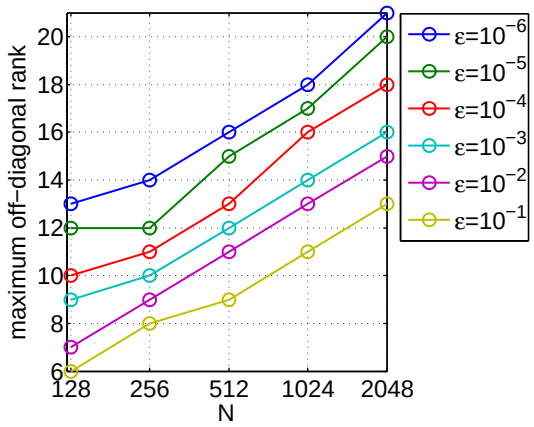

Figure 4-10: Maximum off-diagonal ranks with N for the slow disk, various ϵ. FD error of 10^{-3}, $r_0 = h$.

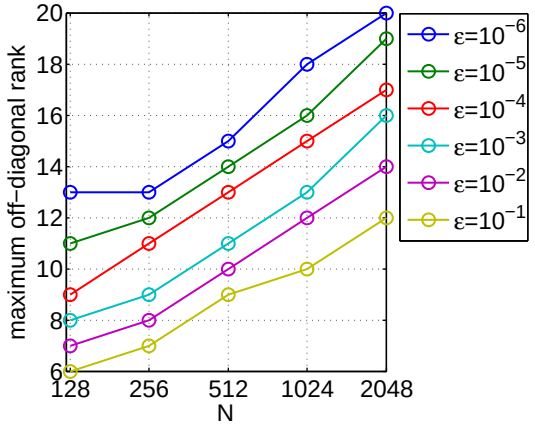

Figure 4-11: Maximum off-diagonal ranks with N for the slow disk, various ϵ. FD error of 10^{-3}, $r_0 = 4h$.

We expect ranks to be slightly smaller for a larger separation r_0 (hence larger C_0) because of Remark 4.3.6. This is indeed the case in our numerical simulations, as we can see by comparing Figures 4-10 ($r_0 = h$) and 4-11 ($r_0 = 4h$), or Figures 4-12 ($r_0 = h$) and 4-13 ($r_0 = 4h$). We can clearly see also how the maximum ranks behave as in the previous theorem, except for the missing square power in J, as alluded to in Remark 4.3.7: they vary logarithmically with k (or N) when the tolerance ϵ is fixed. We expect the slope in a graph of the ranks versus $\log N$ to increase slowly as ϵ becomes smaller, and we see that in fact the slope barely does (from slightly smaller than 2 to slightly larger than 2) as ϵ goes from 10^{-1} to 10^{-6}. Similarly, when we fix N, we expect the ranks to grow logarithmically with $1/\epsilon$, and this is the case. Once again, the slope of the graph with a logarithmic scale for $1/\epsilon$ grows, but only from 1 to 2 or so, as N goes from 128 to 2048.

The off-diagonal ranks of the DtN map for the slow disk behave very similarly as the above for an FD error of 10^{-2}, and also for various other separations r_0. The same is true for the vertical fault, and so we do not show those results.

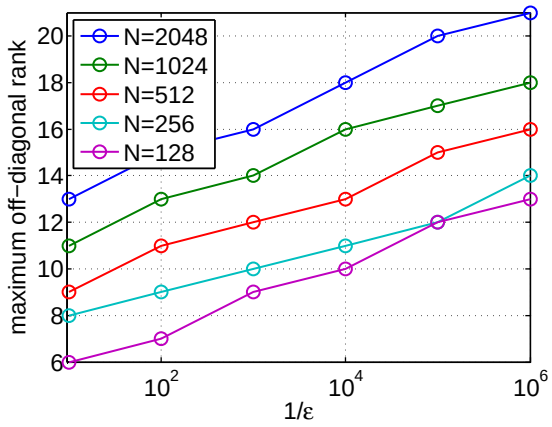

Figure 4-12: Maximum off-diagonal ranks with ε for the slow disk, various N. FD error of 10^{-3}, $r_0 = h$.

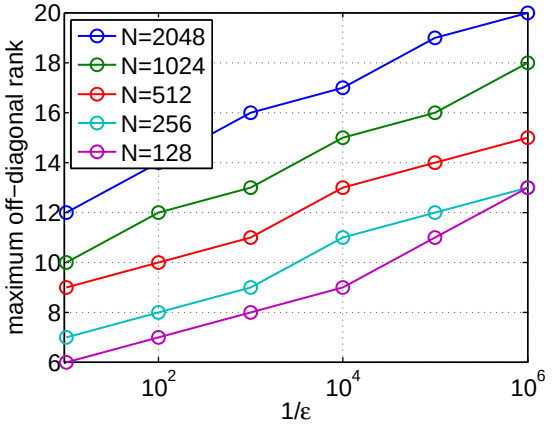

Figure 4-13: Maximum off-diagonal ranks with ε for the slow disk, various N. FD error of 10^{-3}, $r_0 = 4h$.

4.4.2 Constant medium, waveguide, diagonal fault

As for the constant medium, waveguide and diagonal fault, it appears that the term $O(\log^2 k)$ we expect in the size of the ranks is larger than the term $O(\log k |\log \epsilon|)$, especially when the FD error is 10^{-2}. This was mentioned in Remark 4.3.7. As we can see in Figure 4-14 for the diagonal fault, the dependance of the ranks with $\log N$ seems almost quadratic, not linear. This can also be seen in Figure 4-15: here we still see a linear dependence of the ranks with $\log \epsilon$, but we can see that the ranks jump up more and more between different N, as N grows, than they do for the slow disk for example (compare to Figure 4-12). This phenomenon disappears for a smaller FD error (Figures 4-16, 4-17).

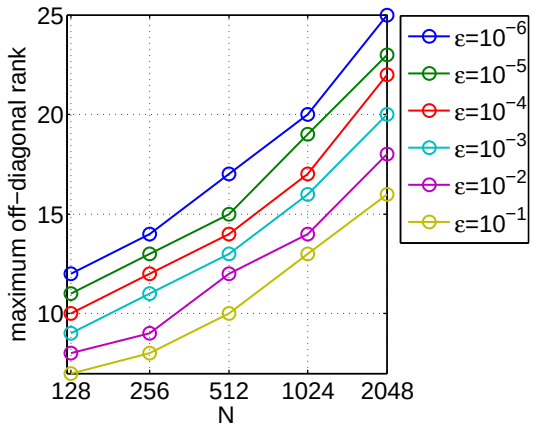

Figure 4-14: Maximum off-diagonal ranks with N for the diagonal fault, various ϵ. FD error of 10^{-2}, $r_0 = h$.

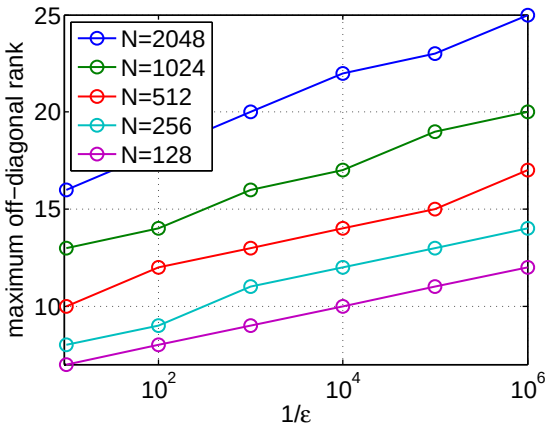

Figure 4-15: Maximum off-diagonal ranks with ε for the diagonal fault, various N. FD error of 10^{-2}, $r_0 = h$.

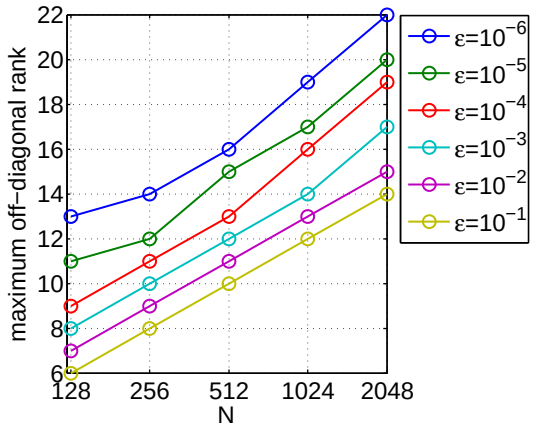

Figure 4-16: Maximum off-diagonal ranks for the diagonal fault as a function of N, various tolerances ϵ. The separation is $r_0 = h$. FD error of 10^{-3}.

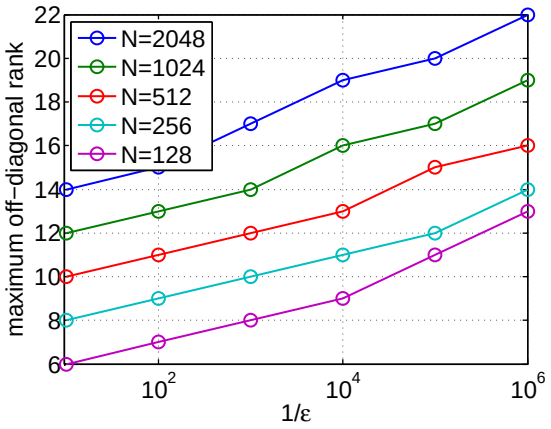

Figure 4-17: Maximum off-diagonal ranks for the diagonal fault as a function of the tolerance ϵ, various N. The separation is $r_0 = h$. FD error of 10^{-3}.

Finally, we also notice that the term $O(\log^2 k)$ seems to remain important compared to the term $O(\log k |\log \epsilon|)$ as the separation r_0 (or C_0) grows, as is somewhat expected from Remark 4.3.6. This can be seen by comparing Figures 4-18 and 4-14, or Figures 4-19 and 4-15.

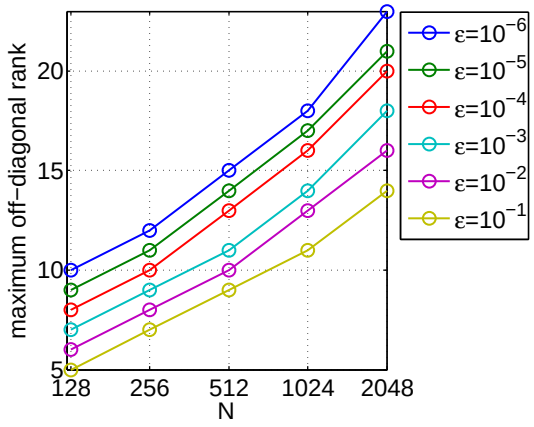

Figure 4-18: Maximum off-diagonal ranks for the diagonal fault as a function of N, various tolerances ϵ. The separation is $r_0 = 8h$. FD error of 10^{-2}.

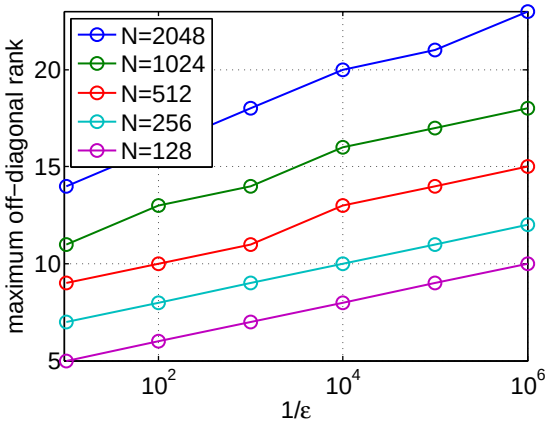

Figure 4-19: Maximum off-diagonal ranks for the diagonal fault as a function of the tolerance ϵ, various N. The separation is $r_0 = 8h$. FD error of 10^{-2}.

4.4.3 Focusing and defocusing media

We have also tested a smooth defocusing medium, that is, a medium where the value of c decreases away from the half-space boundary, tending to a small value far away

from the boundary. The equation we have used is $c(x,y) = 1 + \frac{1}{\pi}\arctan(4(y-1/2))$. This means that as $y \to \infty$, $c \to 3/2$ and that as $y \to -\infty$, $c \to 1/2$. Choosing the half-space to be $y < 0$, we see that c decreases away from $y = 0$ into the negative y's: this is a defocusing medium. We expect the waves in this case to never come back, and so we expect off-diagonal ranks of the DtN map to be very nice, just as in the constant medium case, and this is indeed what happens. We could not see any significant difference between the defocusing medium and the constant medium in terms of off-diagonal ranks.

We have also looked at a focusing medium, that is, one in which c increases away from the interface. This forces waves to come back toward the interface. With the same medium $c(x,y) = 1 + \frac{1}{\pi}\arctan(4(y-1/2))$ as above, but choosing now $y > 1$ as our half-space, we see that c increases away from $y = 1$ into the large positive y's. This is a focusing medium. We have noticed that the off-diagonal ranks of the DtN map for this medium are the same or barely larger than for the constant medium.

This might only mean that the medium we chose did not have many returning waves. A more interesting medium is the following:

$$c(x,y) = 1/2 + |y - 1/2|. \tag{4.21}$$

This linear c has a first derivative bounded away from 0. Of course, this means that solving the Helmholtz equation in this case is much harder, and in particular, the pPML layer needs to be made thicker than for other media. Still, we notice that the ranks are very similar to the previous cases, as we can see in Figures 4-20, 4-21, 4-22, 4-23.

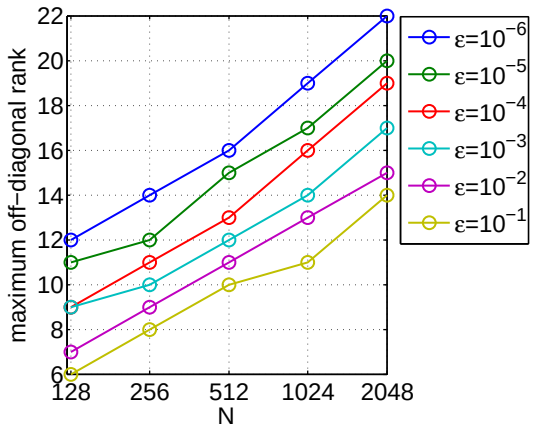

Figure 4-20: Maximum off-diagonal ranks for the focusing medium (4.21) as a function of N, various tolerances ϵ. Separation is $r_0 = h$, FD error of 10^{-2}.

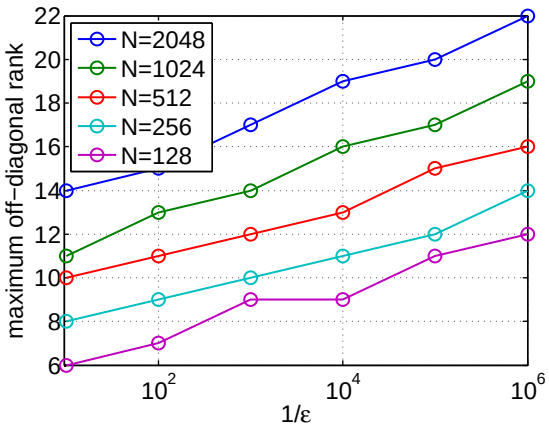

Figure 4-21: Maximum off-diagonal ranks for the focusing medium (4.21) as a function of the tolerance ϵ, various N. Separation is $r_0 = h$, FD error of 10^{-2}.

This could be explained by the fact that we do not go to high enough accuracies to really see those returning waves. Saying this another way, the relative amplitude of the returning waves might be too small to notice at lower accuracies. We were not

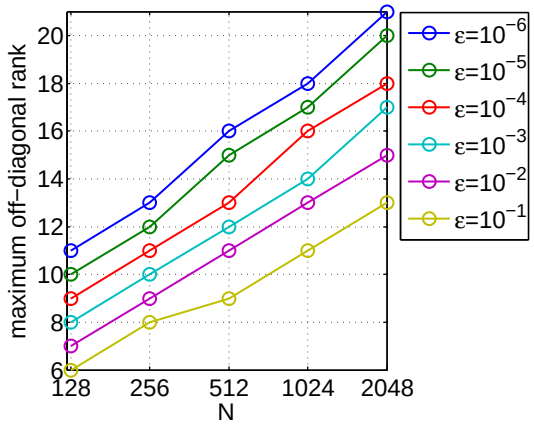

Figure 4-22: Maximum off-diagonal ranks for the focusing medium (4.21) as a function of N, various tolerances ϵ. Separation is $r_0 = h$, FD error of 10^{-3}.

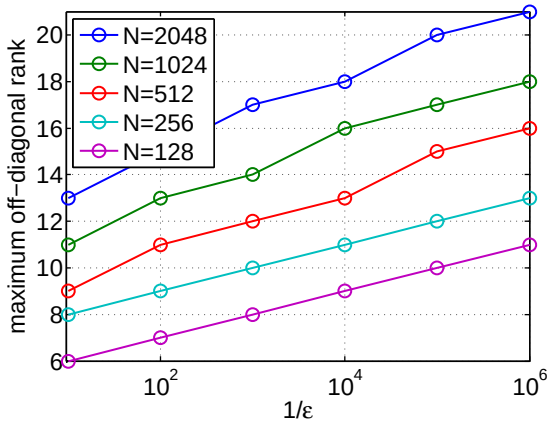

Figure 4-23: Maximum off-diagonal ranks for the focusing medium (4.21) as a function of the tolerance ϵ, various N. Separation is $r_0 = h$, FD error of 10^{-3}.

able to construct an example with returning waves which affected the off-diagonal ranks of the DtN map. Thus we conclude that Theorem 3 holds in a broad range of contexts, at least when ϵ is not very small.

Chapter 5

Conclusion

In this thesis, we have compressed the Dirichlet-to-Neumann map for the Helmholtz equation in two steps, using matrix probing followed by the partitioned low-rank matrix framework. This approach is useful for applications in heterogeneous media because absorbing boundary conditions can be very costly. Especially when one needs to solve the Helmholtz equation with many different sources, it makes sense to perform the precomputation required of our two-step scheme, to speed up subsequent solves.

Probing the DtN map D ultimately makes sense in conjunction with a fast algorithm for its application. In full matrix form, D costs N^2 operations to apply. With the help of a compressed representation, this count becomes p times the application complexity of any atomic basis function B_j, which may or may not be advantageous depending on the particular expansion scheme. The better solution for a fast algorithm, however, is to post-process the compressed expansion from probing into a slightly less compressed, but more algorithmically favorable one, such as hierarchical or partitioned low-rank matrices. These types of matrix structures are not parametrized enough to lend themselves to efficient probing – see for instance [40] for an illustration of the large number of probing vectors required – but give rise to faster algorithms for matrix-vector multiplication. Hence the feasibility of probing, and the availability of a fast algorithm for matrix-vector multiplication, are two different goals that require different expansion schemes.

We found that in order to obtain an efficient compression algorithm, we need to perform some precomputations. The leading cost of those is equivalent to a small number of solves of the original problem, which we can afford if we plan to make many solves in total. Then, a matrix-vector application of the DtN map in a Helmholtz solver is nearly linear in the dimension of the DtN map. The worst-case complexity of a matrix-vector application is in fact super-linear. General oscillatory integrals can often be handled in optimal complexity with the butterfly algorithm [9]. We did not consider this algorithm because currently published methods are not adaptive, nor can they handle diagonal singularities in a kernel-dependent fashion.

Let us summarize the different complexity scalings of our method, recalling that without our compressed ABC, a thousand solves of the free-space problem would cost a thousand times the complexity of one solve, and that complexity depends on

the solver used and on the fact that the ABC used might be very costly. By contrast, the cost of precomputation for matrix probing is of a few (1 to 50) solves of the exterior problem – roughly equivalent to a solve of the free-space problem with the costly ABC. The precomputation cost for PLR compression is $O(NR^2|\mathcal{B}|)$ where $|\mathcal{B}|$ is the number of leaves in the compressed DtN map – empirically $|\mathcal{B}| = O(\log N)$ with a worst case of $|\mathcal{B}| = O(\sqrt{N})$. Finally, the application cost of the compressed ABC is $O(NR|\mathcal{B}|)$, which typically leads to a speed-up of a factor of 30 to 100 for $N \approx 1000$. Using the compressed ABC in a solver, the computational cost of making a thousand solves is then reduced to that of a thousand solves where the ABC is not so costly anymore, but only amounts to a matrix-vector multiplication of nearly linear cost.

Of course, the method presented here has some limitations. The main ones come from the first step of the scheme, matrix probing, which requires a careful design of basis matrices. This is harder to do when the wavelength is comparable to the size of features in the medium, or when the medium has discontinuities.

In addition to the direct applications the proposed two-step scheme has in making a typical (iterative) Helmholtz solver faster when in presence of a heterogeneous media, it could also improve solvers based on Domain Decomposition. Some interesting recent work on fast solvers for the Helmholtz equation has been along the direction of the Domain Decomposition Method (DDM). This method splits up the domain Ω in multiple subdomains. Depending on the method, the subdomains might or might not overlap. In non-overlapping DDM, transmission conditions are used to transfer information about the solution from one domain to the next. Transmission conditions are crucial for the convergence of the algorithm. After all, if there is a source in one subdomain, this will create waves which should travel to other parts of Ω, and this information needs to be transmitted to other subdomains as precisely as possible.

In particular, Stolk in [49] has developed a nearly linear complexity solver for the Helmholtz equation, with transmission conditions based on the PML. Of course, this relies on the PML being an accurate absorbing boundary condition. One may assume that, for faster convergence in heterogeneous media, thicker PMLs might be necessary. In this case, a precomputation to compress the PML might prove useful. Other solvers which rely on transmission conditions are those of Engquist and Ying [20, 19], Zepeda-Núñez et al. [56], Vion and Geuzaine [53].

Another way in which the two-step numerical scheme presented in this thesis could be useful is for compressing the Helmholtz solution operator itself. Indeed, we have compressed ABCs to the DtN map, by using knowledge of the DtN map kernel. If one knows another kernel, one could expand that kernel using matrix probing then compress it with PLR matrices, and obtain a fast application of that solution operator.

Appendix A

Summary of steps

In this appendix, we summarize the various operations needed in each step of the numerical scheme presented in this thesis.

Preliminary remarks

We need two solvers for the numerical scheme, one which does exterior problem solves with boundary condition g on $\partial\Omega$, and one which solves the reformulated problem inside Ω using D or $\overline{D}$ as a boundary condition. Both solvers should be built with the other in mind, so their discretization points agree on $\partial\Omega$.

For the exterior solves, we shall impose a boundary condition g on $\partial\Omega$ (called u_0 in our discussion of layer-stripping), and find the solution on the layer of points just outside of Ω (called u_1 in our discussion of layer-stripping). Note that u_1 has eight more points than u_0 does. However, the four corner points of u_1 are not needed in the normal derivative of u on $\partial\Omega$ (again, because we use the five-point stencil). Also, the four corner points of u_0 need to be used twice each. For example, the solution u_0 at point $(0,0)$ is needed for the normal derivative in the negative x_1 direction (going left) and the normal derivative in the negative x_2 direction (going down). Hence we obtain a DtN operator which takes the $4N$ solution points u_0 (with corners counted twice) to the $4N$ normal derivatives $(u_1 - u_0)/h$ (with corners omitted in u_1). In this way, one can impose a (random) boundary condition on the N points of any one side of $\partial\Omega$ and obtain the resulting Neumann data on the N points of any side of $\partial\Omega$.

Once we have probed and compressed the DtN map D into $\overline{D}$, we will need to use this $\overline{D}$ in a Helmholtz solver. We do this using ghost points u_1 just outside of $\partial\Omega$ (but not at the corners), which we can eliminate using $\overline{D}$. For best results, it is important to use $\overline{D}$ for the same solution points it was obtained. In other words, here we defined $\overline{D}$ as

$$\overline{D}u_0 = \frac{u_1 - u_0}{h}.$$

If instead we use the DtN map as

$$\overline{D}u_1 = \frac{u_1 - u_0}{h} \qquad \text{or} \qquad \overline{D}u_0 = \frac{u_0 - u_{-1}}{h},$$

we lose accuracy. In a nutshell, one needs to be careful about designing the exterior solver and the Helmholtz solver so they agree with each other. We are now ready to consider the steps of the numerical scheme.

A.1 Matrix probing: $D \to \widetilde{D}$

1. Organize information about the various submatrices of D, their multiplicity and the location and "orientation" of each of their copies in D. For example, in the waveguide medium case, the $(1,1)$ submatrix has multiplicity 2, and appears as itself (because of symmetries in c) also in position $(3,3)$. However, the $(2,1)$ submatrix has multiplicity 8 but appears as itself in D only in position $(4,3)$ as well. It appears as its transpose (*not* conjugate transpose) in positions $(1,2)$ and $(3,4)$. It appears with column order flipped in positions $(4,1)$ and $(2,3)$, and finally as its transpose with column order flipped in positions $(1,4)$ and $(3,2)$.

2. Pick a representative for each distinct submatrix of D. To do this, think of which block columns of D will be used for probing. Minimizing the distinct block columns used will minimize the cost of matrix probing by minimizing Q, the sum of all solves needed. See step 4 as well.

3. If the medium c has discontinuities, it might be necessary to split up submatrices further, and to keep track of the "sub-submatrices" and their positions, multiplicities and orientations inside the representative submatrix.

4. Pick a q for each block column, keeping in mind that diagonal submatrices are usually the hardest to probe (hence need a higher p and q), and that submatrices the farthest from the diagonal are typically very easy to probe. It might be wise to pick representatives, in step 2, knowing that some will require a higher q than others.

5. Solve the exterior problem q times on each block column, saving the restriction of the result to the required block rows depending on the representative submatrices you chose. Also, save the random vectors used.

6. For error checking in step 7e, also solve the exterior problem a fixed number of times, say 15, with different random vectors. Again, save both the random vectors and the exterior solves. Use those results to approximate the norm of D as well.

7. For each representative submatrix M (and representative sub-submatrix, if needed), do the following:

   (a) Pick appropriate basis matrices B_j.

   (b) Orthonormalize the basis matrices if necessary. If this step is needed, it is useful to use symmetries in the basis matrices to both reduce the

complexity of the orthonormalization and enforce those symmetries in the orthonormalized basis matrices.

(c) Multiply each basis matrix by the random vectors used in step 5 in solving the exterior problem corresponding to the correct block column of M. Organize results in the matrix $\mathbf{\Psi}$.

(d) Take the pseudo-inverse of $\mathbf{\Psi}$ on the results of the exterior solves from step 5, corresponding to the correct block row of M, to obtain the probing coefficients $\mathbf{c}$ and $\widetilde{M} = \sum c_j B_j$.

(e) To check the probing error, multiply $\widetilde{M}$ with the random vectors used in the exterior solves for error checking purposes, in step 6. Compare to the results of the corresponding exterior solves. Multiply that error by the square root of the multiplicity, and divide by the estimated norm of D.

8. If satisfied with the probing errors of each submatrix, move to next step: PLR compression.

A.2 PLR compression: $\widetilde{D} \to \overline{D}$

1. For each probed representative submatrix $\widetilde{M}$ (and representative sub-submatrix, if needed), do the following:

   (a) Pick a tolerance ε which is smaller than the probing error for that $\widetilde{M}$. A factor of 25 smaller works well usually. Also, pick a maximal desired rank $R_{\max}$. Usually $R_{\max} \leq 8$ for a diagonal submatrix, $R_{\max} \leq 4$ for a submatrix just off of the diagonal, and $R_{\max} = 2$ for a submatrix furthest away from the diagonal work well.

   (b) Compress $\widetilde{M}$ using the PLR compression algorithm. Keep track of the dimensions and ranks of each block, to compare the matrix-vector complexity with that of a dense product.

   (c) Check the error made by the PLR compression by comparing $\widetilde{M}$ and $\overline{M}$, again multiply that error by the square root of the multiplicity, and divide by the estimated norm of D.

2. If satisfied with the PLR errors of each submatrix, move to next step: using the PLR compression of probed submatrices into a Helmholtz solver.

A.3 Using $\overline{D}$ in a Helmholtz solver

Using the Helmholtz solver described in the preliminary remarks of this appendix, obtain the approximate solution $\overline{u}$ from solving with the appropriate boundary condition using $\overline{D}$. Every time a product of a vector v with $\overline{D}$ is needed, build the result from all submatrices of $\overline{D}$. For each submatrix, multiply the correct restriction of that vector v by the correct probed and compressed representative submatrix $\overline{M}$, taking

into account the orientation of the submatrix as discussed in step 1 of the matrix probing part of the numerical scheme.

Bibliography

[1] Milton Abramowitz and Irene A. Stegun. *Handbook of Mathematical Functions*. Dover, New York, 1972.

[2] Daniel Appelö and Tim Colonius. A high-order super-grid-scale absorbing layer and its application to linear hyperbolic systems. *Journal of Computational Physics*, 228(11):4200–4217, 2009.

[3] I. Babuška and S.A. Sauter. Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave number? *SIAM Review*, 42(3):451–484, 2000.

[4] A. Bayliss, C.I. Goldstein, and E. Turkel. On accuracy conditions for the numerical computation of waves. *Journal of Computational Physics*, 59(3):396–404, 1985.

[5] Mario Bebendorf. *Hierarchical Matrices. A Means to Efficiently Solve Elliptic Boundary Value Problems*, volume 63 of *Lecture Notes in Computational Science and Engineering*. Springer, 2008.

[6] Jean-Pierre Bérenger. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of Computational Physics*, 114(2):185–200, 1994.

[7] Brett Borden. Mathematical problems in radar inverse scattering. *Inverse Problems*, 18(1):R1, 2002.

[8] S. Börm, L. Grasedyck, and W. Hackbusch. Hierarchical matrices. Lecture notes 21/2003, Max Planck Institut für Mathematik, Bonn, Germany, Revised 2006.

[9] E. Candès, L. Demanet, and L. Ying. A fast butterfly algorithm for the computation of Fourier integral operators. *Multiscale Modeling and Simulation*, 7(4):1727–1750, 2009.

[10] M. Cheney. A mathematical tutorial on synthetic aperture radar. *SIAM Review*, 43(2):301–312, 2001.

[11] Weng Cho Chew, Jian-Ming Jin, Eric Michielssen, and Jiming Song. *Fast and Efficient Algorithms in Computational Electromagnetics*. Artech House, Boston, 2001.

[12] J. Chiu and L. Demanet. Matrix probing and its conditioning. *SIAM Journal on Numerical Analysis*, 50(1):171–193, 2012.

[13] D. Colton and R. Kress. *Integral Equation Methods in Scattering Theory*. Wiley-Interscience, 1983.

[14] T.A. Davis. Algorithm 832: UMFPACK v4.3—an unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software (TOMS)*, 30(2):196–199, 2011.

[15] L. Demanet, P.D. Letourneau, N. Boumal, H. Calandra, J. Chiu, and S. Snelson. Matrix probing: a randomized preconditioner for the wave-equation hessian. *Appl. Comput. Harmon. Anal.*, 32:155–168, 2012.

[16] L. Demanet and L. Ying. Scattering in flatland: Efficient representations via wave atoms. *Foundations of Computational Mathematics*, 10(5):569–613, 2010.

[17] J. Diaz and P. Joly. A time domain analysis of PML models in acoustics. *Comput. Methods Appl. Mech. Engrg.*, 195:3820–3853, 2006.

[18] B. Engquist and A. Majda. Absorbing boundary conditions for the numerical simulation of waves. *Mathematics of Computation*, 31(139):629–651, 1977.

[19] B. Engquist and L. Ying. Sweeping preconditioner for the Helmholtz equation: Moving perfectly matched layers. *Multiscale Modeling And Simulation*, 9(2):686–710, 2011.

[20] Bjorn Engquist and Lexing Ying. Sweeping preconditioner for the Helmholtz equation: Hierarchical matrix representation. *Communications on Pure and Applied Mathematics*, 64(5):697–735, 2011.

[21] Y. Erlangga, C. Oosterlee, and C. Vuik. A novel multigrid based preconditioner for heterogeneous Helmholtz problems. *SIAM Journal on Scientific Computing*, 27(4):1471–1492, 2006.

[22] Yogi A. Erlangga. Advances in iterative methods and preconditioners for the Helmholtz equation. *Archives of Computational Methods in Engineering*, 15(1):37–66, 2008.

[23] Sonia Fliss and Patrick Joly. Wave propagation in locally perturbed periodic media (case with absorption): Numerical aspects. *Journal of Computational Physics*, 231(4):1244 – 1271, 2012.

[24] G. B. Folland. *Introduction to Partial Differential Equations*. Princeton University Press, Princeton, NJ, second edition, 1995.

[25] A. Gillman, A. Barnett, and P.G. Martinsson. A spectrally accurate direct solution technique for frequency-domain scattering problems with variable media. *In review.*

[26] A. Gillman and P.G. Martinsson. A direct solver with O(N) complexity for variable coefficient elliptic PDEs discretized via a high-order composite spectral collocation method. *In review.*

[27] L. Grasedyck and W. Hackbusch. Construction and arithmetics of H-matrices. *Computing*, 70(4):295–334, 2003.

[28] Robert E. Green, Paul McIntire, and Albert S Birks. *Ultrasonic testing.* American Society for Nondestructive Testing, Columbus, OH, 2nd edition, 1991.

[29] W. Hackbusch. A sparse matrix arithmetic based on H-matrices. i. introduction to H-matrices. *Computing*, 62(2):89–108, 1999.

[30] W. Hackbusch, B. N. Khoromskij, and R. Kriemann. Hierarchical matrices based on a weak admissibility criterion. *Computing*, 73(3):207–243, 2004.

[31] T. Hagstrom and T. Warburton. Complete radiation boundary conditions: minimizing the long time error growth of local methods. Technical report, Southern Methodist University.

[32] N. Halko, P. Martinsson, and J. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.

[33] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade. *Photonic Crystals: Molding the Flow of Light.* Princeton University Press, second edition, 2008.

[34] P. Jones, J. Ma, and V. Rokhlin. A fast direct algorithm for the solution of the Laplace equation on regions with fractal boundaries. *Journal of Computational Physics*, 113(1):35–51, 1994.

[35] B. R. Julian, A. Ross, G. R. Foulger, and J. R. Evans. Three-dimensional seismic image of a geothermal reservoir: The Geysers, California. *Geophysical Research Letters*, 23(6):685–688, 1996.

[36] H. Keller and M. Lentini. Invariant imbedding, the box scheme and an equivalence between them. *SIAM Journal on Numerical Analysis*, 19(5):942–962, 1982.

[37] Joseph B. Keller and Dan Givoli. Exact non-reflecting boundary conditions. *Journal of Computational Physics*, 82(1):172–192, 1989.

[38] L. F. Knockaert and D. De Zutter. On the complex symmetry of the Poincaré-Steklov operator. *Progress In Electromagnetics Research B*, 7:145–157, 2008.

[39] Edo Liberty, Franco Woolfe, Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. Randomized algorithms for the low-rank approximation of matrices. *Proceedings of the National Academy of Sciences*, 104(51):20167–20172, 2007.

[40] Lin Lin, Jianfeng Lu, and Lexing Ying. Fast construction of hierarchical matrix representation from matrix-vector multiplication. *Journal of Computational Physics*, 230(10):4071–4087, 2011.

[41] P.-R. Loh, A. F. Oskooi, M. Ibanescu, M. Skorobogatiy, and S. G. Johnson. Fundamental relation between phase and group velocity, and application to the failure of perfectly matched layers in backward-wave structures. *Physical Review E*, 79:065601, 2009.

[42] P.G. Martinsson and V. Rokhlin. A fast direct solver for scattering problems involving elongated structures. *Journal of Computational Physics*, 221(1):288–302, 2007.

[43] W. McLean. *Strongly Elliptic Systems and Boundary Integral Equations*. Cambridge University Press, 2000.

[44] A. F. Oskooi, L. Zhang, Y. Avniel, and S. G. Johnson. The failure of perfectly matched layers, and their redemption by adiabatic absorbers. *Optics Express*, 16:11376–11392, 2008.

[45] Stefan A. Sauter and Christoph Schwab. *Boundary element methods*, volume 39 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, 2011. Translated and expanded from the 2004 German original.

[46] J. A. Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93(4):1591–1595, 1996.

[47] E. Somersalo, M. Cheney, D. Isaacson, and E. Isaacson. Layer stripping: a direct numerical method for impedance imaging. *Inverse Problems*, 7(6):899, 1991.

[48] Josef Stoer and Roland Bulirsch. *Introduction to Numerical Analysis*. Springer, third edition, 2002.

[49] Christiaan C. Stolk. A rapidly converging domain decomposition method for the Helmholtz equation. *Journal of Computational Physics*, 241(0):240–252, 2013.

[50] E. Tadmor. The exponential accuracy of Fourier and Chebyshev differencing methods. *SIAM Journal on Numerical Analysis*, 23(1):1–10, 1986.

[51] Lloyd N. Trefethen and David Bau, III. *Numerical linear algebra*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997.

[52] Jayaram K. Udupa and Gabor T. Herman. *3D Imaging in Medicine, Second Edition*. CRC Press, 1999.

[53] A. Vion and C. Geuzaine. Double sweep preconditioner for optimized Schwarz methods applied to the Helmholtz problem. *Journal of Computational Physics*, 266(0):171–190, 2014.

[54] Öz Yilmaz. *Seismic Data Analysis*. Society of Exploration Geophysicists, 2001.

[55] Jianhua Yuan, Ya Yan Lu, and Xavier Antoine. Modeling photonic crystals by boundary integral equations and Dirichlet-to-Neumann maps. *Journal of Computational Physics*, 227(9):4617 – 4629, 2008.

[56] L. Zepeda-Núñez, R. J. Hewett, and L. Demanet. Preconditioning the 2D Helmholtz equation with polarized traces. *Preprint*, 2014.